



US011914869B2

(12) **United States Patent**  
**Dain et al.**

(10) **Patent No.:** **US 11,914,869 B2**  
(45) **Date of Patent:** **Feb. 27, 2024**

(54) **METHODS AND SYSTEMS FOR ENCRYPTION BASED ON INTELLIGENT DATA CLASSIFICATION**

G06F 3/0607; G06F 3/067; G06F 12/1408; G06F 12/1458; G06F 2212/1052; G06F 2212/402

See application file for complete search history.

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

(72) Inventors: **Joseph Dain**, Vail, AZ (US); **Nilesh P. Bhosale**, Pune (IN); **Abhishek Jain**, Uttar Pradesh (IN); **Gregory Kishi**, Oro Valley, AZ (US)

5,953,419 A 9/1999 Lohstroh et al.  
7,269,729 B2 9/2007 He et al.  
(Continued)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

FOREIGN PATENT DOCUMENTS

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1011 days.

EP 2223237 B1 8/2016

OTHER PUBLICATIONS

(21) Appl. No.: **16/257,450**

Ghane, A.H. et al. "A network steganographic approach to overlay cognitive radio systems utilizing systematic coding," Physical Communication, vol. 27, 2018.

(22) Filed: **Jan. 25, 2019**

(Continued)

(65) **Prior Publication Data**

*Primary Examiner* — Baboucarr Faal

US 2020/0241769 A1 Jul. 30, 2020

(74) *Attorney, Agent, or Firm* — Scully, Scott, Murphy & Presser, P.C.

(51) **Int. Cl.**  
**G06F 3/06** (2006.01)  
**G06F 12/14** (2006.01)

(Continued)

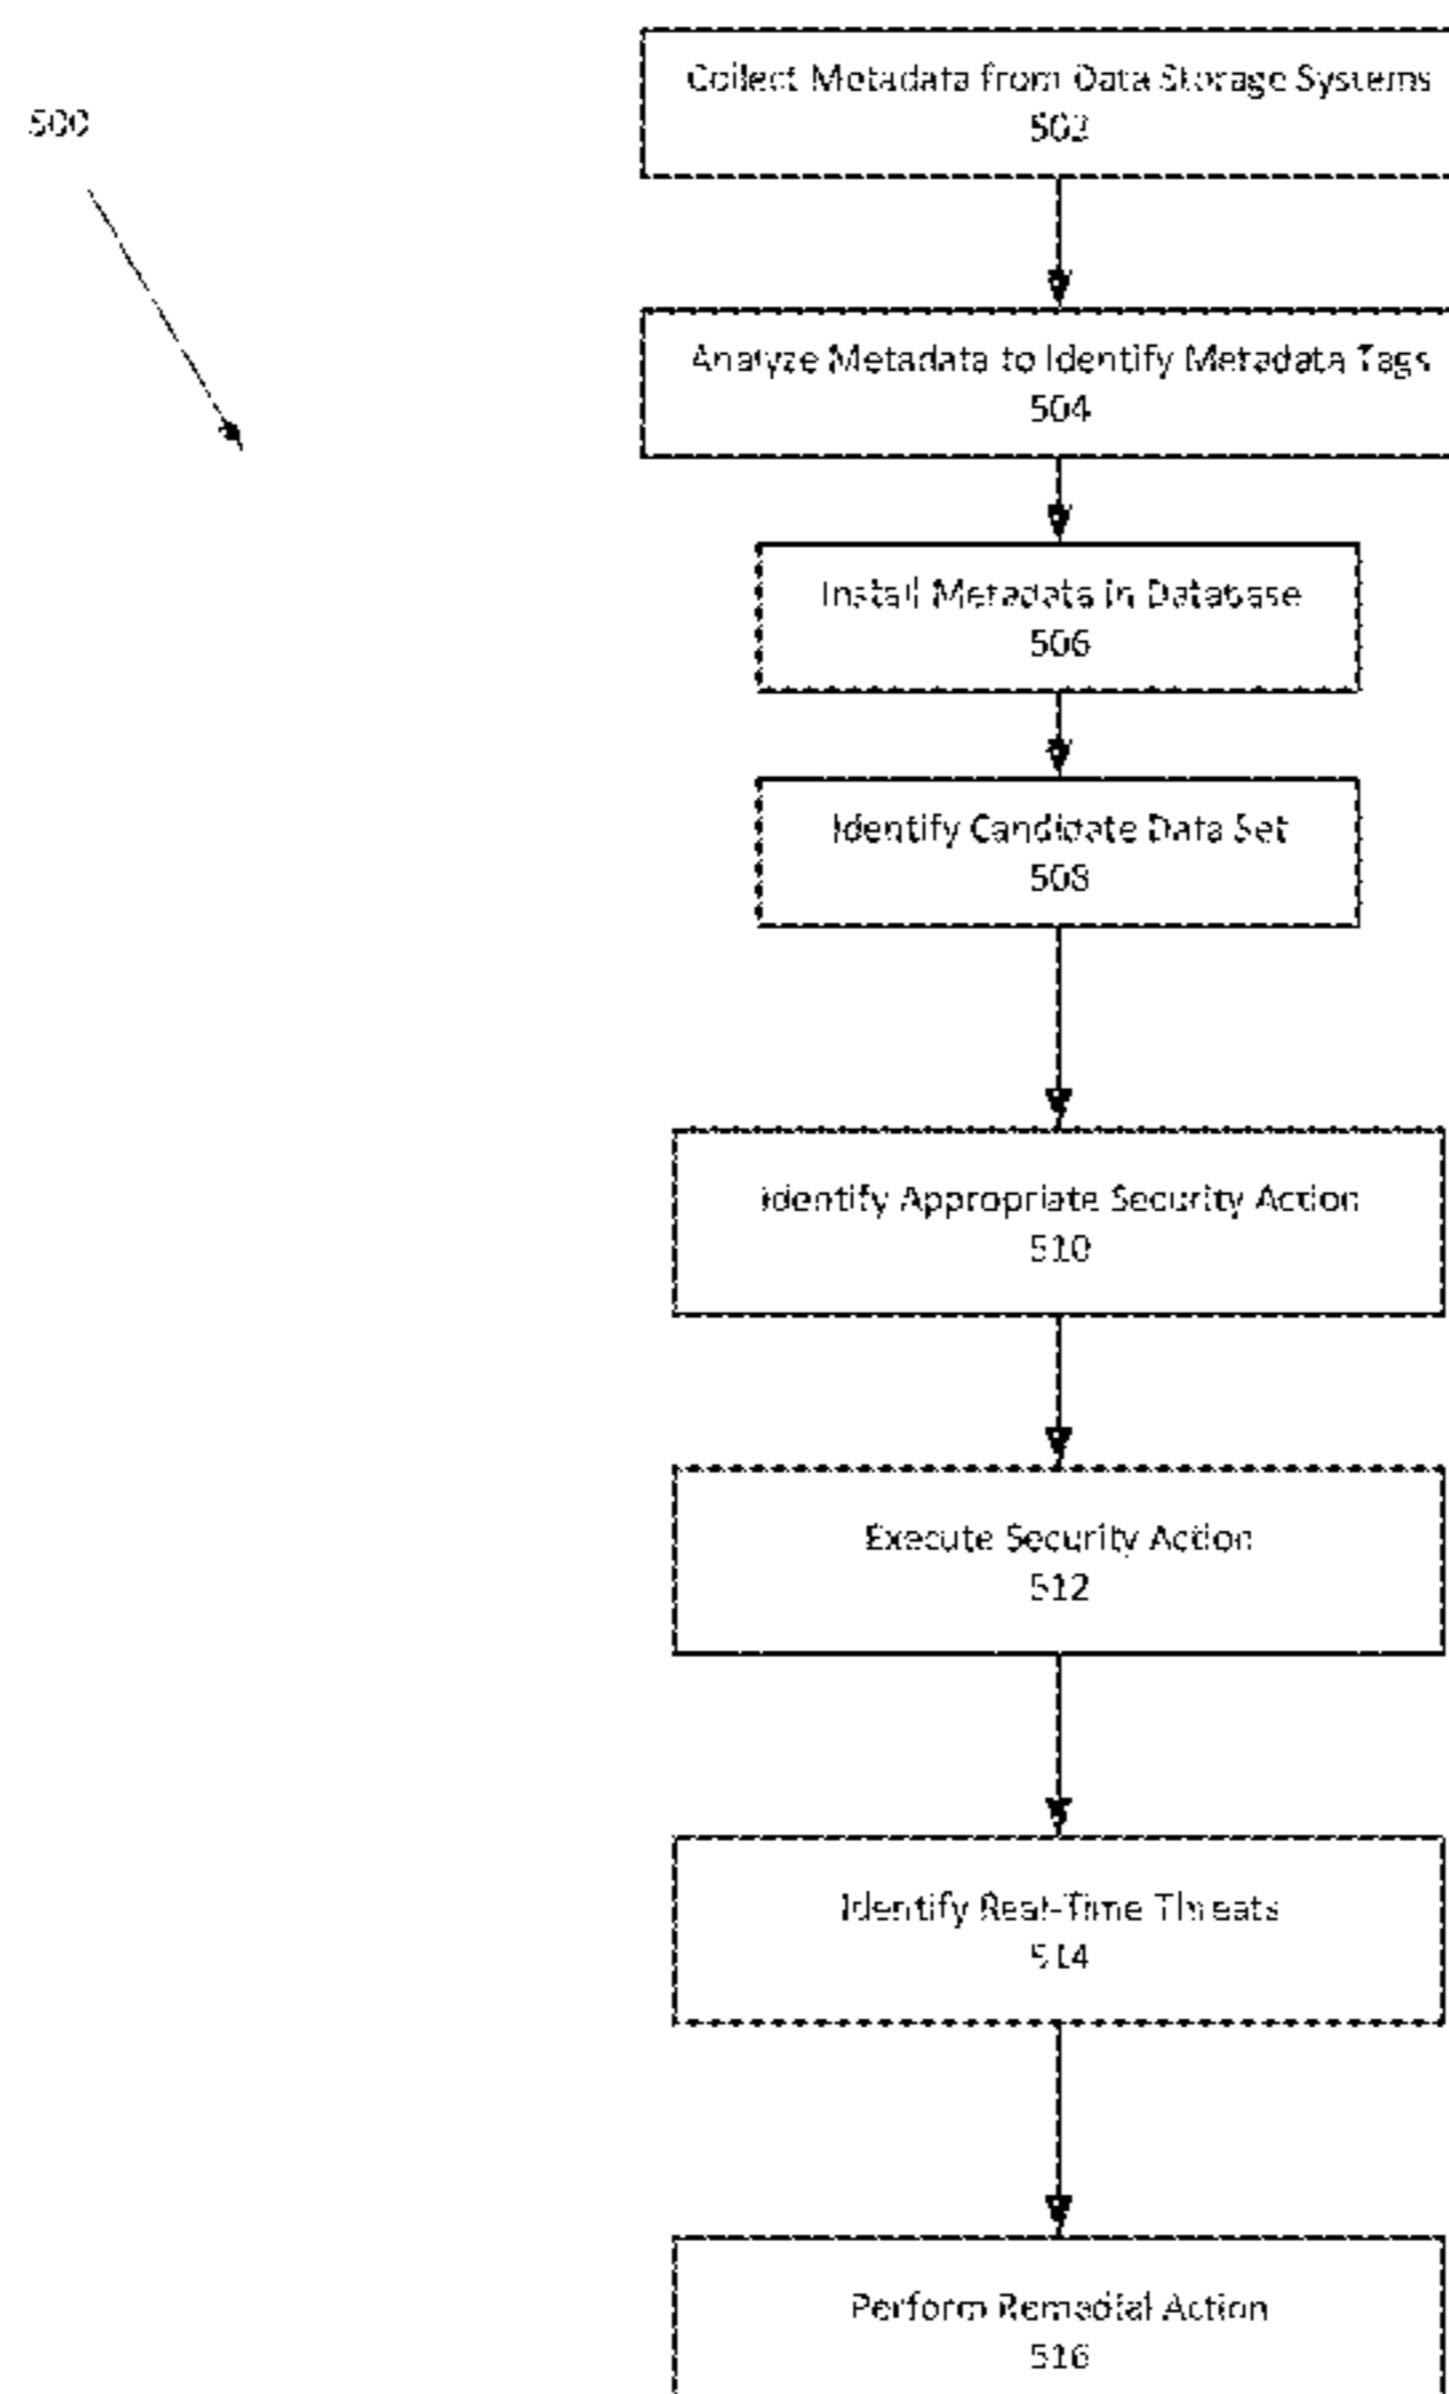
(57) **ABSTRACT**

(52) **U.S. Cl.**  
CPC ..... **G06F 3/0623** (2013.01); **G06F 3/067** (2013.01); **G06F 3/0607** (2013.01); **G06F 12/1408** (2013.01); **G06F 12/1458** (2013.01); **G06F 16/258** (2019.01); **G06F 16/907** (2019.01); **G06F 2212/1052** (2013.01); **G06F 2212/402** (2013.01)

Systems and methods for cognitive encryption of data are disclosed. The methods may include maintaining a plurality of data storage systems in communication with an external metadata management system, operating the metadata management system to store metadata corresponding to data residing on the plurality of data storage systems, identifying a candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed using information included in the metadata management system, and in response to identifying the candidate data set, identifying the at least one security action.

(58) **Field of Classification Search**  
CPC .... G06F 3/0623; G06F 16/907; G06F 16/258;

**18 Claims, 5 Drawing Sheets**



- (51) **Int. Cl.**  
**G06F 16/25** (2019.01)  
**G06F 16/907** (2019.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,610,285	B1	10/2009	Zoellner et al.	
7,752,314	B2	7/2010	Trevor et al.	
7,849,048	B2	12/2010	Angseth et al.	
8,195,630	B2	6/2012	Ellis et al.	
8,224,788	B2	7/2012	Schachter	
8,407,216	B2	3/2013	Walker et al.	
8,489,650	B2	7/2013	Li et al.	
8,689,239	B2	4/2014	Taylor et al.	
8,769,055	B2	7/2014	Murphy et al.	
8,898,101	B2	11/2014	Bone et al.	
8,935,366	B2	1/2015	Mehr et al.	
9,002,906	B1	4/2015	Rice et al.	
9,130,844	B1	9/2015	Malik et al.	
9,507,807	B1	11/2016	Florissi et al.	
9,571,666	B2	2/2017	Sharma	
9,665,404	B2	5/2017	Hu et al.	
9,720,965	B1	8/2017	Miskie	
10,152,384	B1*	12/2018	Amit .....	G06F 21/6227
10,951,705	B1	3/2021	Desai et al.	
2007/0038822	A1*	2/2007	Correl .....	G06F 11/1466 711/162
2009/0083940	A1	4/2009	Albrecht et al.	
2009/0089340	A1	4/2009	Ohr et al.	
2009/0119576	A1*	5/2009	Pepper .....	G06F 40/169 707/999.102
2010/0274750	A1	10/2010	Oltean et al.	
2011/0167277	A1*	7/2011	Kitamura .....	G06K 15/4095 358/1.14
2011/0304774	A1	12/2011	Latta et al.	
2011/0307493	A1	12/2011	Noor	
2012/0304247	A1	11/2012	Badger et al.	
2013/0298050	A1	11/2013	Emanuel et al.	
2015/0112939	A1	4/2015	Cantwell et al.	
2015/0154509	A1	6/2015	Lightner et al.	
2015/0242263	A1	8/2015	Klose	
2015/0334203	A1	11/2015	Soundararajan et al.	
2016/0004734	A1	1/2016	Naryzhny et al.	
2016/0019317	A1	1/2016	Pawar et al.	
2016/0063021	A1	3/2016	Morgan et al.	
2016/0063993	A1	3/2016	Dolan et al.	
2016/0110370	A1	4/2016	Wetherall et al.	
2016/0110374	A1	4/2016	Wetherall et al.	
2016/0170844	A1	6/2016	Long et al.	
2016/0196324	A1	7/2016	Haviv et al.	
2016/0267086	A1	9/2016	Liden	
2016/0371500	A1*	12/2016	Huang .....	G06F 21/6218
2017/0005923	A1	1/2017	Babakian	
2017/0161520	A1	6/2017	Lockhart, III et al.	
2017/0169103	A1	6/2017	Juneja et al.	
2017/0177904	A1	6/2017	Bilodeau et al.	
2017/0179995	A1	6/2017	Oviedo	
2017/0220487	A1*	8/2017	Jung .....	G06F 12/1408
2017/0235764	A1	8/2017	Sharpe et al.	
2017/0337393	A1	11/2017	Crowe et al.	
2018/0075166	A1	3/2018	Pawar et al.	
2018/0225177	A1	8/2018	Bhagi et al.	
2019/0018870	A1	1/2019	Bhagwat	
2019/0190921	A1*	6/2019	Rieser .....	H04L 63/105
2019/0325059	A1	10/2019	Fair	
2020/0026799	A1	1/2020	Fukuda et al.	
2020/0285737	A1*	9/2020	Kraus .....	G06F 21/552

OTHER PUBLICATIONS

Ogiela, L. et al., "Classification of cognitive systems dedicated to data sharing," Optics and Photonics for Information Processing XI, vol. 10395, International Society for Optics and Photonics, 2017.  
Hacigümüs, H. et al. "Providing Database as a Service," Data Engineering, 2002.

Mattsson, U.T., "Database Encryption—How to Balance Security with Performance," Protegrity Corp., 2005.  
Siersdorfer, S. et al., "Automatic Video Tagging Using Content Redundancy," SIGIR '09 Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval, Jul. 19-23, 2009, pp. 395-402.  
Dawei, Y. et al., "A Probabilistic Model for Personalized Tag Prediction," KDD '10 Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Jul. 25-28, 2010, pp. 959-968.  
Lipczak, M. et al., "Learning in Efficient Tag Recommendation," RecSys '10 Proceedings of the Fourth ACM Conference on Recommender Systems, Sep. 26-30, 2010, pp. 167-174.  
"Storage Tiering with Dell EMC ISILON SmartPools" Dell, Inc., Nov. 2017.  
Next Generation Storage Tiering with EMC ISILON SmartPools, EMC Corporation, Apr. 2013.  
Disclosed Anonymously: Minimizing the Costly Update of Persistent Meta-Data in Multi-Tier Environment by Using Metadata Filter, ip.com, Nov. 24, 2015.  
Disclosed Anonymously: Method for Notification Events Classification Based on Tags, ip.com, May 13, 2011.  
Wang, Y. et al., "Gnothi: Separating Data and Metadata for Efficient and Available Storage Replication," USENIX ATC, 2012.  
Custom Metadata Types Implementation Guide, Salesforce, Aug. 9, 2017.  
IBM: "A Tag Library Based Approach for Auto-Refreshing Web Pages Based on External Events", ip.com, Feb. 6, 2007.  
Exner, P. et al., "KOSHIK—A Large-scale Distributed Computing Framework for NLP," ICPRAM, 2014.  
Niekler, A. et al., "Leipzig Corpus Miner—A Text Mining Infrastructure for Qualitative Data Analysis," arXiv preprint, 2017.  
Doyle, A. et al., "The EMBERS Architecture for Streaming Predictive Analytics," 2014 IEEE International Conference on Big Data, 2014.  
Osazee-Obazee, E.P., "Data Management in an Elastic Storage Environment," Helsinki Metropolia University of Applied Sciences, May 30, 2015.  
Jensen, K.B., "High Throughput Tertiary Storage in HPC Environments," Proceedings of the Doctoral Symposium of the 17th International Middleware Conference. ACM, 2016.  
Calhoun, S.P. et al., "Large scale research data archiving: Training for an inconvenient technology," Journal of Computational Science, 2016.  
Reuther, A. et al., "LLSuperCloud: Sharing HPC Systems for Diverse Rapid Prototyping," High Performance Extreme Computing Conference (HPEC), IEEE, 2013.  
Tan, J. et al., "DynMR: Dynamic MapReduce with ReduceTask Interleaving and MapTask Backfilling," Proceedings of the Ninth European Conference on Computer Systems. ACM, 2014.  
Li, C.S. et al., "Software defined environments: An introduction," IBM Journal of Research and Development, vol. 58, No. 2/3, Mar./May 2014.  
Disclosed Anonymously, Data aware optimized backups—Object & System level, Jun. 2, 2016.  
Disclosed Anonymously, System and Method to implement smart database restore techniques for heterogeneous databases to enhance the repair algorithm using meta data, Feb. 18, 2016.  
Akolkar, A. et al., Intelligent Virtual Machine Backups Using Events and Alarms of Virtual Infrastructure, Jun. 17, 2011.  
Holzenkamp, C., Capture, Indexing & Auto-Categorization, AIIM International, 2002.  
List of IBM Patents or Applications Treated as Related.  
Lovelace, et al., "IBM Tivoli Storage Manager as a Data Protection Solution", IBM, 2014, p. 7 and pp. 36-50.  
Office Action dated Apr. 3, 2020 received in U.S. Appl. No. 16/257,498, 36 pages.  
Office Action dated Apr. 9, 2020 received in U.S. Appl. No. 16/257,425, 22 pages.  
Office Action dated Apr. 14, 2020 received in U.S. Appl. No. 16/257,434, 24 pages.  
Office Action dated Oct. 20, 2021 received in U.S. Appl. No. 16/257,457.

(56)

**References Cited**

OTHER PUBLICATIONS

Office Action dated Mar. 24, 2022 received in U.S. Appl. No. 16/257,457.

Office Action dated Apr. 16, 2020 received in U.S. Appl. No. 16/257,489 50 pages.

Office Action dated Sep. 3, 2021 received in U.S. Appl. No. 16/257,457.

Office Action dated May 14, 2021 received in U.S. Appl. No. 16/257,457.

Office Action dated Feb. 17, 2021 received in U.S. Appl. No. 16/257,476.

Office Action dated Sep. 19, 2022 received in U.S. Appl. No. 16/257,457.

Office Action dated Aug. 24, 2023 received in a related U.S. Patent Application, namely U.S. Appl. No. 16/257,457.

Office Action dated Dec. 7, 2023 received in U.S. Appl. No. 16/257,457.

\* cited by examiner

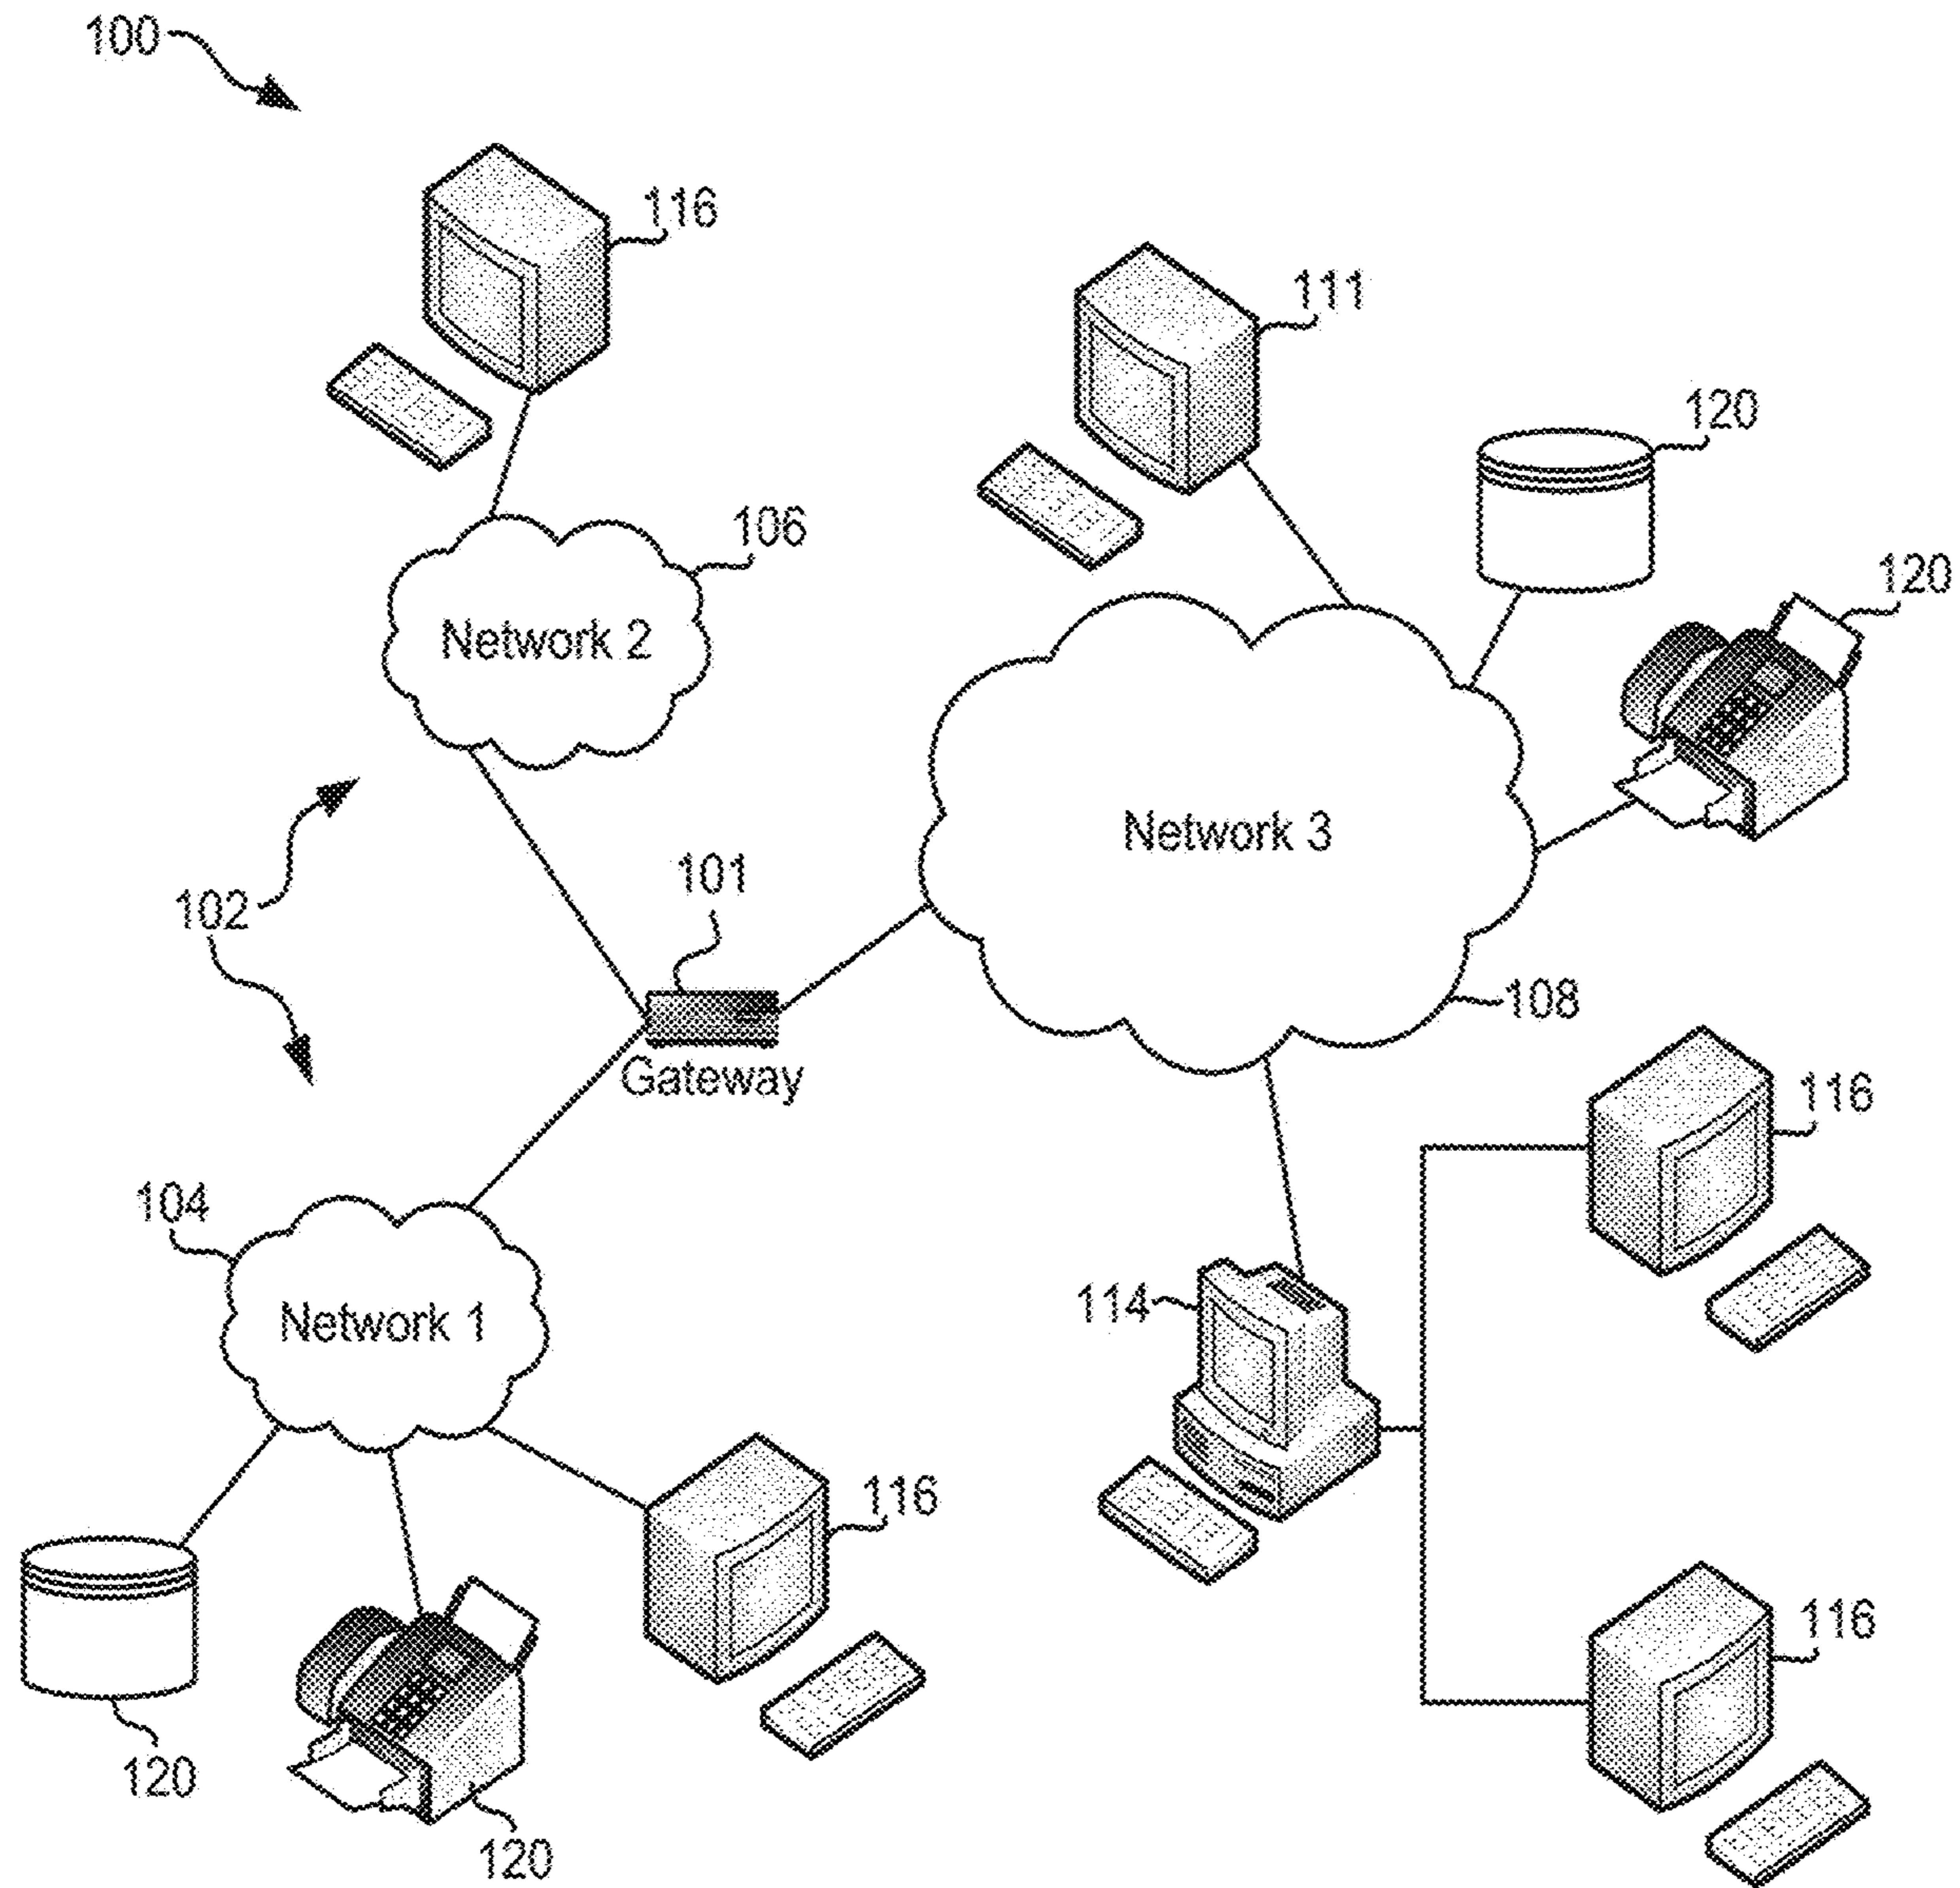


FIG. 1

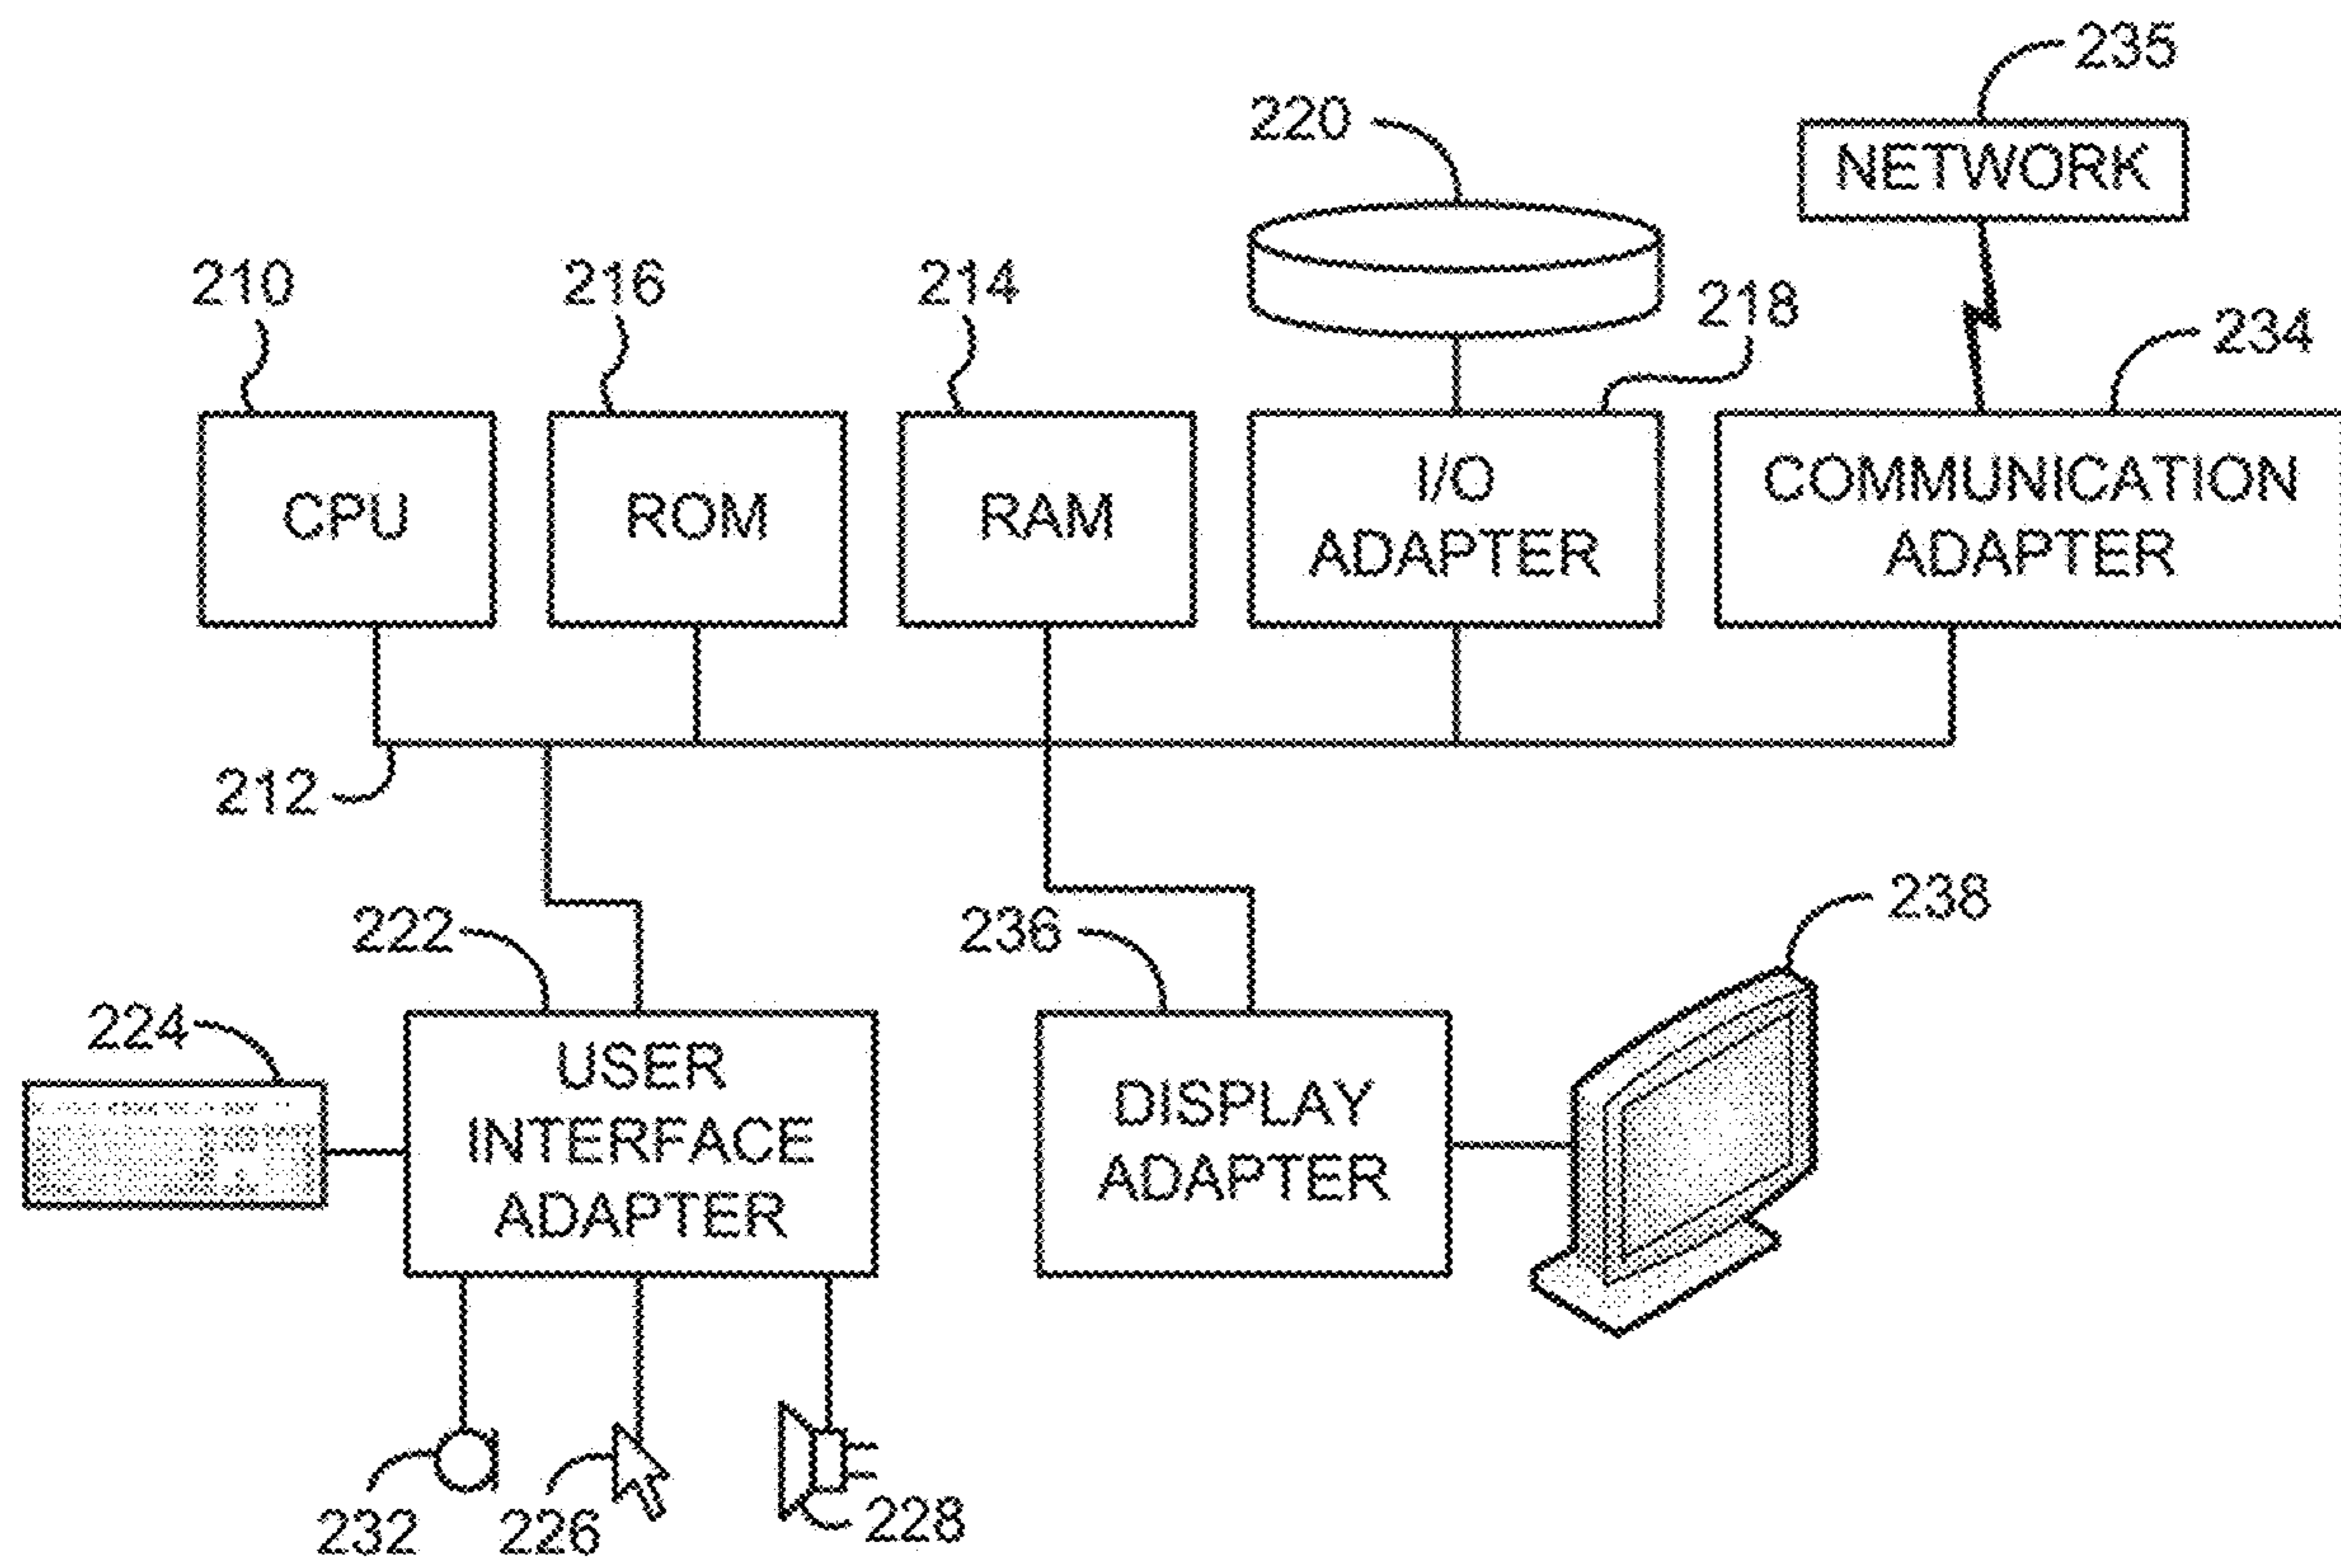


FIG. 2

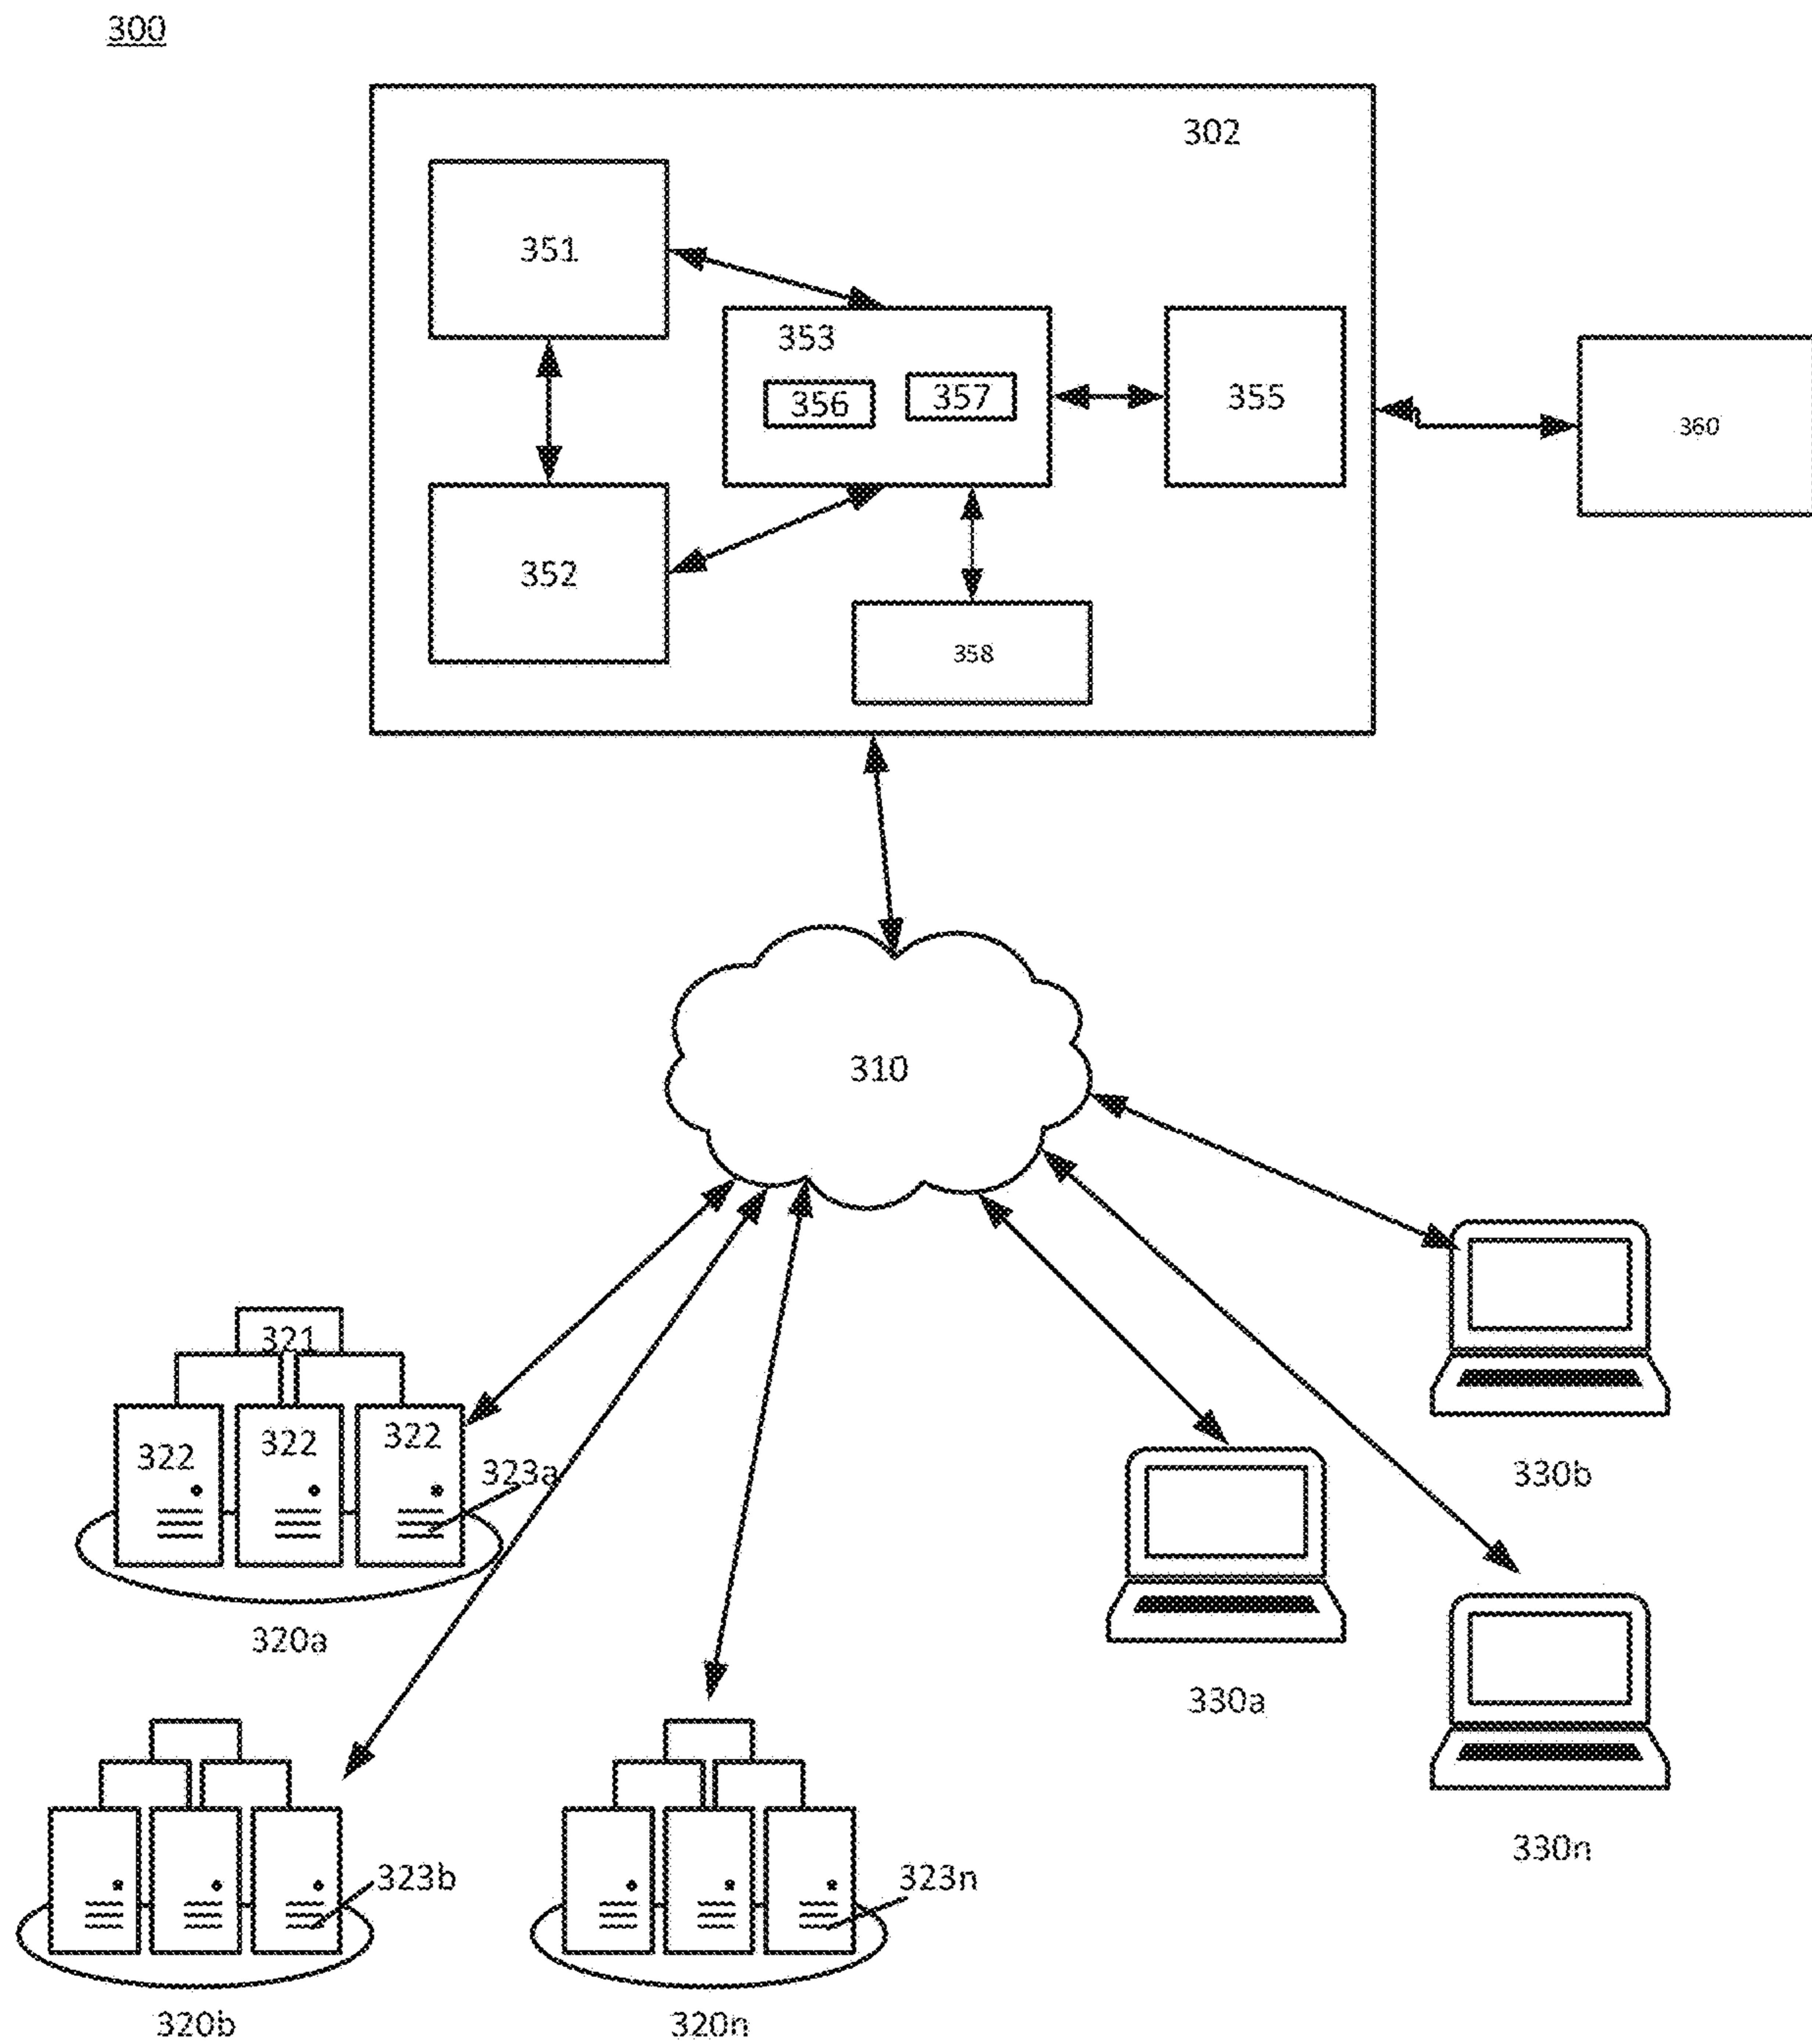


FIG. 3

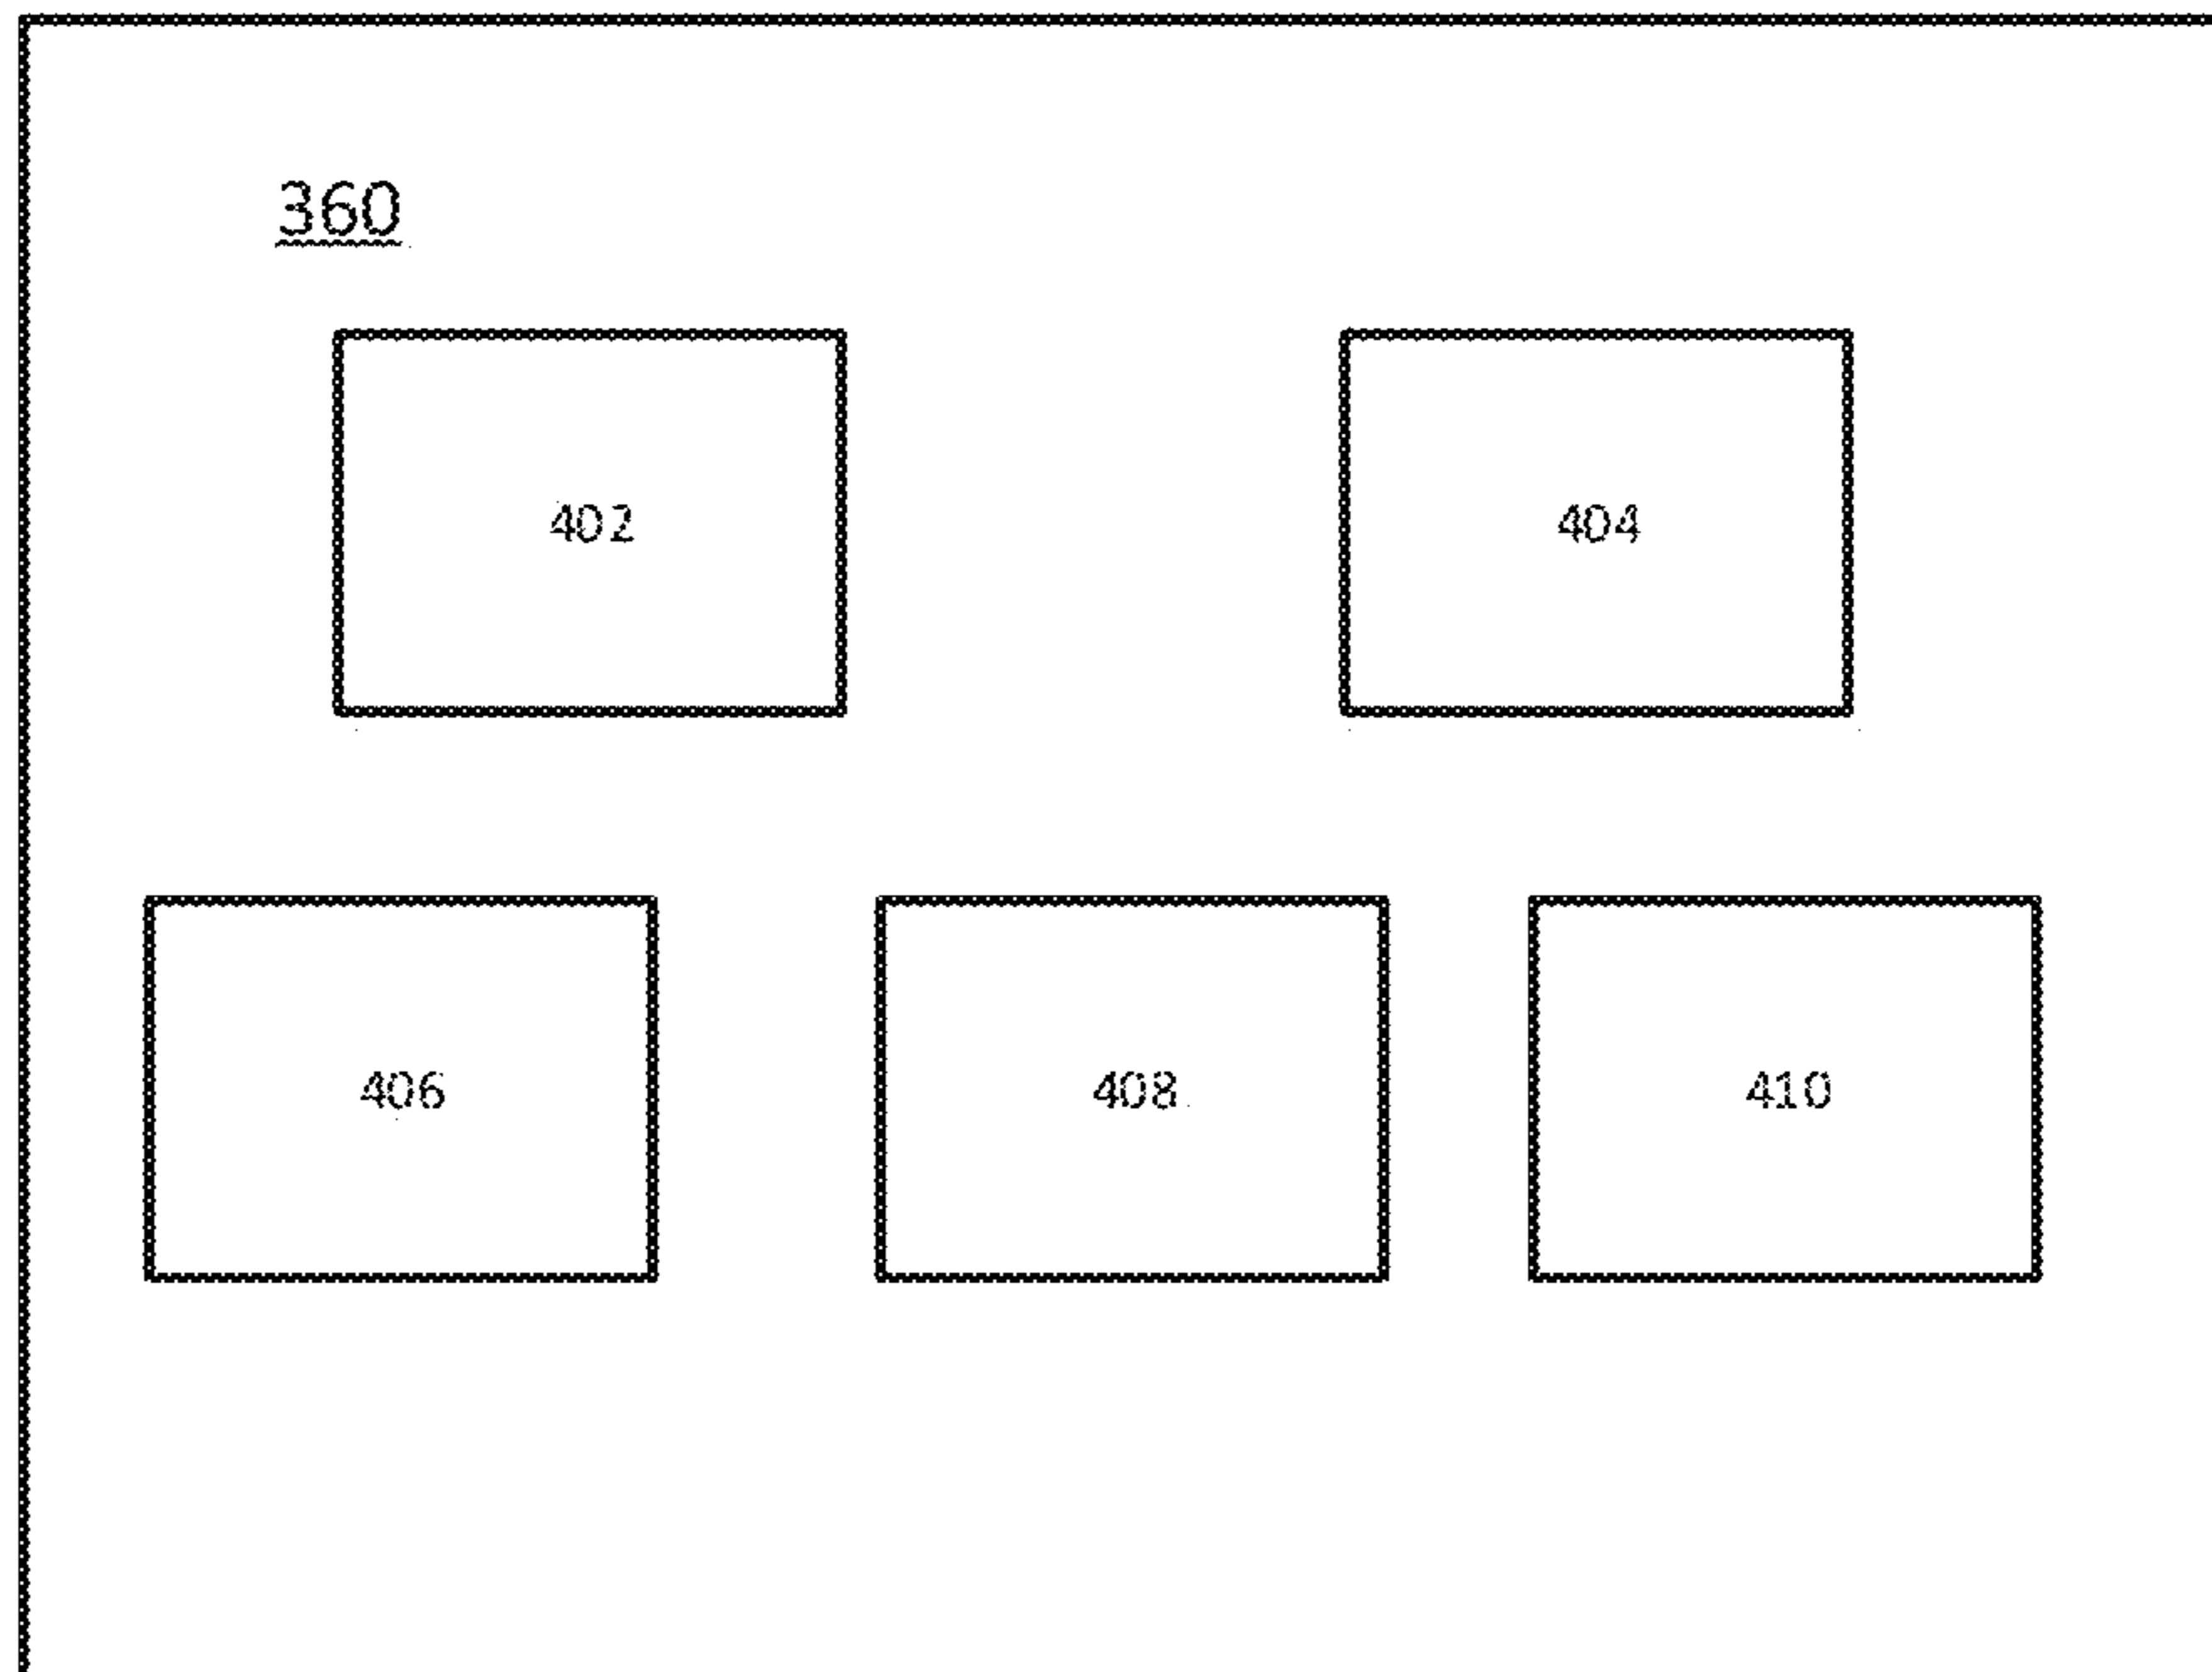


FIG. 4

500  
↘

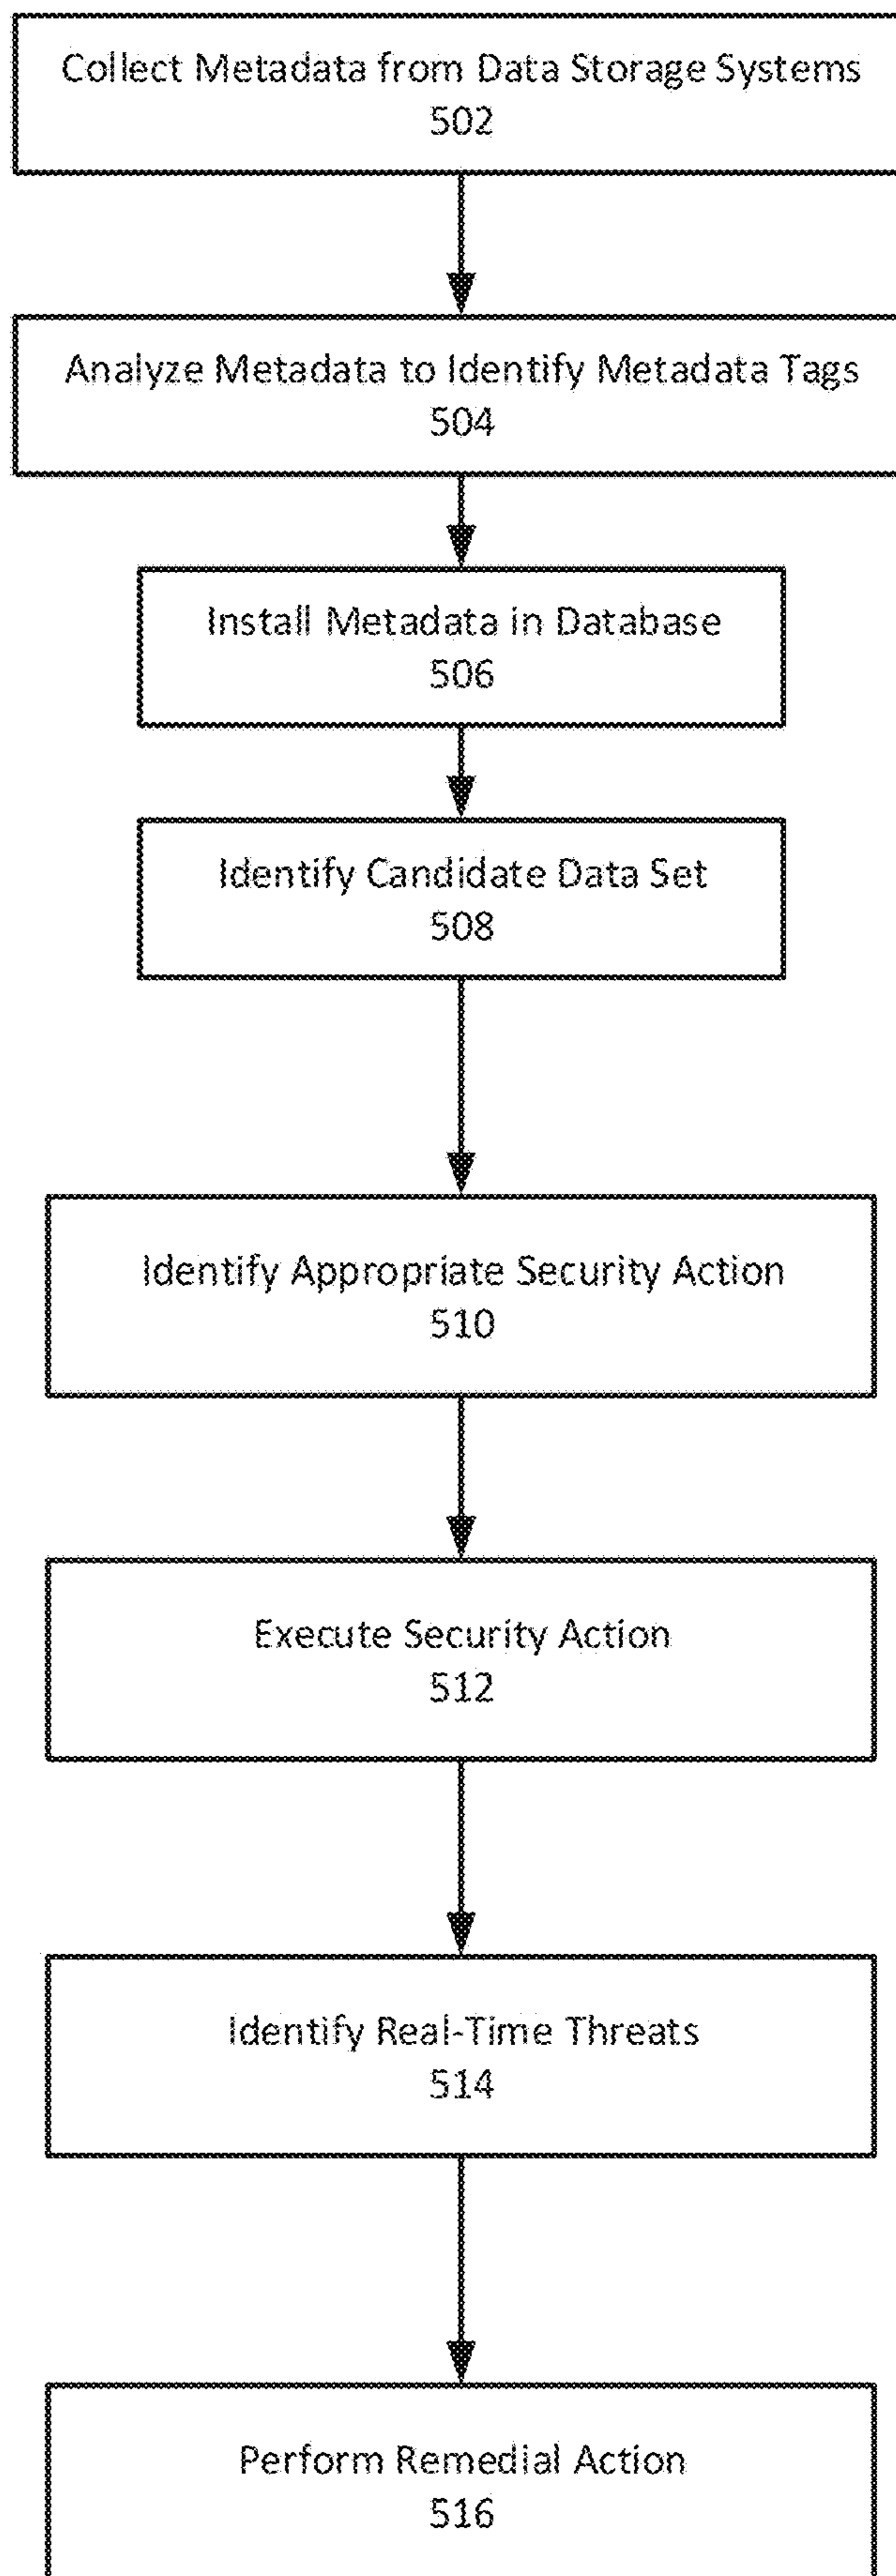


FIG. 5



**METHODS AND SYSTEMS FOR  
ENCRYPTION BASED ON INTELLIGENT  
DATA CLASSIFICATION**

BACKGROUND

The disclosure herein relates generally to a dynamic data protection methods and systems. Specifically, the present disclosure relates to an intelligent system for identifying and encrypting sensitive data in a storage system.

Users and organizations that deal with significant quantities of digital information often have difficulty managing files and data in an efficient and intuitive manner. An inability to easily store, organize, and locate documents and content, while causing difficulty and irritation at the level of the individual user, may translate into significant inefficiencies and lost opportunities at the organizational level. Lost documents, overlooked e-mails and records, and the duplication of work between users or departments may impact a business's productivity and agility. For the digital consumer, difficulty organizing and locating digital data may result in user frustration and the accidental re-purchasing of extant content.

Modern high-capacity hard drives and remote storage solutions allow for the retention of large numbers of documents and records nearly indefinitely; however, increases in storage capacity have often not been accompanied by a corresponding increase in the effectiveness of document management tools and technology. Most modern storage solutions utilize some combination of a traditional directory-based file system and search-based data management such as full-text search or basic keyword tagging. Although appropriate for some types of data, both types of systems may present significant challenges when dealing with large numbers of files or heterogeneous data sets. Directory-based solutions may be satisfactory for highly structured data or content; however, directory-trees often break down as an organizational method when a document or datum is relevant across one or more data categories or when a user desires to cross-reference or locate documents based on an alternate organizational schema. Simple text and keyword search-based systems generally discard the rigid structure of the directory-tree, but may present other challenges, such as requiring that the user remember specific terms or phrases associated with the document to be located. The lack of structure associated with many keyword or full-text based data management solutions may also pose difficulties when similar keyword terms occur over different classes of documents, such as a "flight" keyword being used both for trip records and engineering documents.

Some of the weaknesses with directory and keyword/text search-based systems may be mitigated by associating metadata with each piece of data. Metadata is broadly defined as "data about data" i.e. a label or description. Thus, a given item of metadata may be used to describe an individual datum, or a content item, and/or a collection of data which can include a plurality of content items. The fundamental role of metadata is to facilitate or aid in the understanding, use and management of data. The metadata required for efficient data management is dependent on, and varies with, the type of data and the context of use of this data. Using as an example a library, the data is the content of the titles stocked, and the metadata about a title would typically include a description of the content, and any other information relevant for whatever purposes, for example the publication date, author, location in the library, etc. For photographic images, metadata typically labels the date the

photograph was taken, whether day or evening, the camera settings, and information related to copyright control, such as the name of the photographer, and owner and date of copyright. Therefore, the metadata may have clearer semantics and include some category information to organize the data in the repository. Even more, the relationships among different metadata items may be involved to describe more complex semantics. Obviously, the query on metadata is more effective to retrieve appropriate results than the full-text search, especially for some specific areas difficult to apply the full-text search, such as multimedia.

The content stored on data storage systems also includes sensitive and personal information. The rapid increase in the amount of data stored in data storage systems also leads to challenges relating to adequate protection of such sensitive information and personal information included in the stored data. Determining and enforcing the correct level of encryption and/or access control for classified content is typically a manual process subject to human error. While metadata may be used to identify the correct level of encryption and/or access control, with the increase in the amount and complexity of metadata, effective searching of metadata also becomes difficult.

Furthermore, existing methods fail to identify new items of sensitive information or personal information, new loopholes in encryption methods used, or the like.

SUMMARY

The summary of the disclosure is given to aid understanding of a data storage system, data storage system architectural structure, processor, and method of encrypting contents of a data storage system, and not with an intent to limit the disclosure or the invention. The present disclosure is directed to a person of ordinary skill in the art. It should be understood that various aspects and features of the disclosure may advantageously be used separately in some instances, or in combination with other aspects and features of the disclosure in other instances. Accordingly, variations and modifications may be made to the computer system, the architectural structure, processor, and their method of operation to achieve different effects.

According to an embodiment of the present disclosure, systems and methods for cognitive encryption of data are disclosed. The methods may include maintaining a plurality of data storage systems in communication with an external metadata management system, operating the metadata management system to store metadata corresponding to data residing on the plurality of data storage systems, identifying a candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed using information included in the metadata management system, and in response to identifying the candidate data set, identifying the at least one security action.

In one or more embodiments, the methods may also include executing the at least one security action on the candidate data set. Optionally, the at least one security action may include encryption of the candidate data set.

In certain embodiments, identifying the at least one security action may include extracting one or more facets of the candidate data set stored with the metadata in the metadata management system, and using the one or more facets to identify the at least one security action. Optionally, the extracted one or more facets may be identified by performing data analytics on the candidate data set. Alternatively and/or additionally, the extracted one or more facets may be iden-

ified by performing data analytics on at least one component of metadata corresponding to the candidate data set.

In certain embodiments, identifying the at least one security action may include identifying one or more custom tags for metadata corresponding to the candidate data set, and using the one or more custom tags to identify the at least one security action.

In at least one embodiment, identifying the candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed may include receiving a query from a user that includes one or more rules for selecting the candidate data set using metadata stored in the metadata management system. Optionally, identifying the candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed may include identifying the candidate data set based on metadata received in response to a data operation event performed on the candidate data set.

In one or more embodiments, the methods may also include identifying a security level associated with the candidate data set based on at least one of the group consisting of: one or more facets extracted from the candidate data set and stored with the metadata in the metadata management system, one or more facets extracted from metadata associated with the candidate data set and stored with the metadata in the metadata management system, and/or one or more custom tags corresponding to metadata associated with the candidate data set. The methods may also include using the security level to identify the at least one security action.

In some embodiments, the methods may also include receiving a real-time alert comprising a threat to the at least one security action, and identifying at least one remedial action for countering the threat.

The above methods may be performed by a processor executing instructions stored on non-transitory computer readable medium.

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular descriptions of exemplary embodiments of the invention as illustrated in the accompanying drawings wherein like reference numbers generally represent like parts of exemplary embodiments of the invention.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The various aspects, features and embodiments of a computer system, computer architectural structure, processor, and their method of operation will be better understood when read in conjunction with the figures provided. Embodiments are provided in the figures for the purpose of illustrating aspects, features, and/or various embodiments of the computer system, computer architectural structure, processors, and their method of operation, but the claims should not be limited to the precise arrangement, structures, features, aspects, assemblies, sub assemblies, systems, circuitry, embodiments, or devices shown, and the arrangements, structures, subassemblies, assemblies, features, aspects, methods, processes, circuitry, embodiments, and devices shown may be used singularly or in combination with other arrangements, structures, assemblies, subassemblies, systems, features, aspects, circuitry, embodiments, methods and devices.

FIG. 1 depicts one example of a computing environment, according to embodiments of the present disclosure.

FIG. 2 is a functional block diagram illustrating a computer system, according to embodiments of the present disclosure.

FIG. 3 depicts an example block diagram of an information management system, according to embodiments of the present disclosure.

FIG. 4 is a functional block diagram illustrating a security module, according to embodiments of the present disclosure.

FIG. 5 is an exemplary flowchart illustrating and describing a method for intelligent encryption of data in data storage systems according to embodiments of the present disclosure.

#### DETAILED DESCRIPTION

The following description is made for illustrating the general principles of the invention and is not meant to limit the inventive concepts claimed herein. In the following detailed description, numerous details are set forth in order to provide an understanding of the computer system, computer architectural structure, storage systems, processor, and their method of operation, however, it will be understood by those skilled in the art that different and numerous embodiments of the computer system, computer architectural structure, storage systems, processor, and their methods of operation may be practiced without those specific details, and the claims and disclosure should not be limited to the embodiments, assemblies subassemblies, assemblies, processes, methods, aspects, or details specifically described and shown herein. Further, particular features described herein can be used in combination with other described features in each of the various possible combinations and permutations.

Unless otherwise specifically defined herein, all terms are to be given their broadest possible interpretation including meanings implied from the specification as well as meanings understood by those skilled in the art and/or as defined in dictionaries, treatises, etc. It must also be noted that, as used in the specification and the appended claims, the singular forms “a,” “an” and “the” include plural referents unless otherwise specified, and that the terms “comprises” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

As used herein, the term “content” or “data” means any computer-readable data including, but not limited to, digital photographs, digitized analog photos, music files, video clips, text documents, interactive programs, web pages, word processing documents, computer assisted design files, blueprints, flowcharts, invoices, database reports, database records, video game assets, sound samples, transaction log files, electronic documents, files which simply name other objects, and the like. The content may be organized and stored in the form of objects, files, blocks, or any other suitable format in one or more data storage systems, and can include files, directories, file system volumes, data blocks, extents, or any other hierarchies or organizations of data blocks. As used herein, a “data set” can refer to (1) any file that is currently addressable by a file system or that was previously addressable by the file system (e.g., an archive file) and/or (2) a subset of such a file (e.g., a data block). Data may include structured data (e.g., database files), unstructured data (e.g., documents), and/or semi-structured data. Specifically, a data set can be a file, directory, share, volume, region within a volume, or an embedded object. Data sets can be complex, containing other embedded

objects. For example, a file can be a container containing other files, or a volume can have a file system on top of it which in turn contains files. The system is capable of recognizing complex objects and tracking changes at finer embedded object granularity.

A “container” may be a data set which may have other embedded objects, such as a file, directory, file system, or volume.

As used herein, the term “metadata” refers to any descriptive or identifying information in computer-processable form that is associated with particular content or data set. Generally speaking, content will have metadata that is relevant to a number of characteristics of the content and the overall content collection (e.g., a file), including, but not limited to, the content’s technical aspects (format, bytes used, date of creation), the workflow in which the content participates (creator, owner, publisher, date of publication, copyright information, etc) and the subject matter of the content (the nature of the sound of an audio file, be it music or a sound-effect, the subject of a photograph or video clip, the abstract of a lengthy text document, excerpted particulars of invoices or other data-interchange format files). For example, metadata items may include but are not limited to one or more of the following: the content owner (e.g., the client or user that generates the content), the last modified time (e.g., the time of the most recent modification of a data set), a data set name (e.g., a file name), a data set size (e.g., a number of bytes of data set), information about the content (e.g., an indication as to the existence of a particular search term), user-supplied tags, to/from information for email (e.g., an email sender, recipient, etc.), creation date, file type (e.g., format or application type), last accessed time, application type (e.g., type of application that generated the data block), location/network (e.g., a current, past or future location of the data set and network pathways to/from the data block), geographic location (e.g., GPS coordinates), frequency of change (e.g., a period in which the data set is modified), business unit (e.g., a group or department that generates, manages or is otherwise associated with the set), aging information (e.g., a schedule, such as a time period, in which the data set is migrated to secondary or long term storage), boot sectors, partition layouts, file location within a file folder directory structure, user permissions, owners, groups, access control lists [ACLS]), system metadata (e.g., registry information), combinations of the same or the other similar information related to the data set.

The term “metadata tag” or “tag” refers to any descriptive or identifying information in computer-processable form that is associated with particular metadata, and that is indicative of the actual information of the content included in various data storage systems and with which the metadata is associated.

The following discussion omits or only briefly describes conventional features of data storage systems and information processing systems, including processors and micro-processor systems and architectures, which are apparent to those skilled in the art. It is assumed that those skilled in the art are familiar with the general architecture of data storage system, and in particular operations of data storage systems for data storage and/or operations on stored data. It may be noted that a numbered element is numbered according to the figure in which the element is introduced, and is typically referred to by that number throughout succeeding figures.

FIG. 1 illustrates an architecture 100, in accordance with one embodiment. As shown in FIG. 1, a plurality of remote networks 102 are provided including a first remote network 104 and a second remote network 106. A gateway 101 may

be coupled between the remote networks 102 and a proximate network 108. In the context of the present architecture 100, the networks 104, 106 may each take any form including, but not limited to a LAN, a WAN such as the Internet, public switched telephone network (PSTN), internal telephone network, etc.

In use, the gateway 101 serves as an entrance point from the remote networks 102 to the proximate network 108. As such, the gateway 101 may function as a router, which is capable of directing a given packet of data that arrives at the gateway 101, and a switch, which furnishes the actual path in and out of the gateway 101 for a given packet.

Further included is at least one data server 114 coupled to the proximate network 108, and which is accessible from the remote networks 102 via the gateway 101. It should be noted that the data server(s) 114 may include any type of computing device/groupware. Coupled to each data server 114 is a plurality of user devices 116. User devices 116 may also be connected directly through one of the networks 104, 106, 108. Such user devices 116 may include a desktop computer, lap-top computer, hand-held computer, printer, smartphone, or any other type of logic device. It should be noted that a user device 111 may also be directly coupled to any of the networks, in one embodiment.

A peripheral 120 or series of peripherals 120, e.g., facsimile machines, printers, networked and/or local storage units or systems, etc., may be coupled to one or more of the networks 104, 106, 108. It should be noted that databases and/or additional components may be utilized with, or integrated into, any type of network element coupled to the networks 104, 106, 108. In the context of the present description, a network element may refer to any component of a network.

According to some approaches, methods and systems described herein may be implemented with and/or on virtual systems and/or systems which emulate one or more other systems, such as a UNIX system which emulates an IBM z/OS environment, a UNIX system which virtually hosts a MICROSOFT WINDOWS environment, a MICROSOFT WINDOWS system which emulates an IBM z/OS environment, etc. This virtualization and/or emulation may be enhanced through the use of VMWARE software, in some embodiments.

In more approaches, one or more networks 104, 106, 108, may represent a cluster of systems commonly referred to as a “cloud.” In cloud computing, shared resources, such as processing power, peripherals, software, data, servers, etc., are provided to any system in the cloud in an on-demand relationship, thereby allowing access and distribution of services across many computing systems. Cloud computing typically involves an Internet connection between the systems operating in the cloud, but other techniques of connecting the systems may also be used.

FIG. 2 shows a representative hardware environment associated with a user device 116 and/or server 114 of FIG. 1, in accordance with one embodiment. Such figure illustrates a typical hardware configuration of a workstation having a central processing unit 210, such as a microprocessor, and a number of other units interconnected via a system bus 212.

The workstation shown in FIG. 2 includes a Random Access Memory (RAM) 214, Read Only Memory (ROM) 216, an I/O adapter 218 for connecting peripheral devices such as disk storage units 220 to the bus 212, a user interface adapter 222 for connecting a keyboard 224, a mouse 226, a speaker 228, a microphone 232, and/or other user interface devices such as a touch screen and a digital camera (not

shown) to the bus **212**, communication adapter **234** for connecting the workstation to a communication network **235** (e.g., a data processing network) and a display adapter **236** for connecting the bus **212** to a display device **238**.

The workstation may have resident thereon an operating system such as Microsoft Windows® Operating System (OS), MAC OS, UNIX OS, etc. It will be appreciated that a preferred embodiment may also be implemented on platforms and operating systems other than those mentioned. A preferred embodiment may be written using XML, C, and/or C++ language, or other programming languages, along with an object oriented programming methodology. Object oriented programming (OOP), which has become increasingly used to develop complex applications, may be used.

Referring now to FIG. 3, there is illustrated an example block diagram of an information management system **300** that includes a set of networked data storage systems **320a**, **320b** . . . **320n**, client devices **330a**, **330b** . . . **330n**, and a metadata management system **302** in communication via a data network **310** and in accordance with implementations of this disclosure. It can be appreciated that the implementations disclosed herein are not limited by the number of storage devices or data storage systems attached to data network **310**. It can be further appreciated that storage devices or data storage systems attached to data network **310** are not limited by communication protocols, storage environment, physical location, etc.

In one embodiment, each data storage system **320a**, **320b** . . . **320n** may include a storage subsystem **321** and storage devices **322**. The storage subsystem **321** may comprise a storage server or an enterprise storage server, such as the IBMS Enterprise Storage Server®. (IBM and Enterprise Storage Server are registered trademarks of IBM). The storage devices **322** may comprise storage systems known in the art, such as a Direct Access Storage Device (DASD), Just a Bunch of Disks (JBOD), a Redundant Array of Independent Disks (RAID), a virtualization device, tape storage, optical disk storage, or any other storage system known in the art. The storage devices **322** may comprise content organized as object storage, file storage, and/or block storage. In certain embodiments, multiple storage subsystems may be implemented in one storage subsystem **321** and storage devices **322**, or one storage subsystem may be implemented with one or more storage subsystems having attached storage devices.

In certain embodiments, client devices **330a**, **330b** . . . **330n** may be general purpose computers having a plurality of components. These components may include a central processing unit (CPU), main memory, I/O devices, and storage devices (for example, flash memory, hard drives and others). The main memory may be coupled to the CPU via a system bus or a local memory bus. The main memory may be used to provide the CPU access to data and/or program information that is stored in main memory at execution time. Typically, the main memory is composed of random access memory (RAM) circuits. A computer system with the CPU and main memory is often referred to as a host system. The client devices **330a**, **330b** . . . **330n** can have at least one operating system (e.g., Microsoft Windows, Mac OS X, iOS, IBM z/OS, Linux, other Unix-based operating systems, etc.) installed thereon, which may support or host one or more file systems and other applications.

The data storage systems **320a**, **320b** . . . **320n** and client devices **330a**, **330b** . . . **330n** communicate according to well-known protocols, such as the Network File System (NFS) or the Common Internet File System (CIFS) protocols, to make content stored on data storage systems **320a**,

**320b** . . . **320n** appear to users and/or application programs as though the content were stored locally on the client systems **330a**, **330b** . . . **330n**. In a typical mode of operation, the client devices **330a**, **330b** . . . **330n** transmit one or more input/output commands, such as an NFS or CIFS request, over the computer network **310** to the data storage systems **320a**, **320b** . . . **320n**, which in turn issues an NFS or CIFS response containing the requested content over the network **310** to the respective client devices **330a**, **330b** . . . **330n**.

The client devices **330a**, **330b** . . . **330n** may execute (internally and/or externally) one or more applications, which generate and manipulate the content on the one or more data storage systems **320a**, **320b** . . . **320n**. The applications generally facilitate the operations of an organization (or multiple affiliated organizations), and can include, without limitation, mail server applications (e.g., Microsoft Exchange Server), file server applications, mail client applications (e.g., Microsoft Exchange Client), database applications (e.g., SQL, Oracle, SAP, Lotus Notes Database), word processing applications (e.g., Microsoft Word), spreadsheet applications, financial applications, presentation applications, browser applications, mobile applications, entertainment applications, and so on. The applications may also have the ability to access (e.g., read and write to) data storage systems **320a**, **320b** . . . **320n** using a network file system protocol such as NFS or CIFS.

As shown, the data storage systems **320a**, **320b** . . . **320n**, the client devices **330a**, **330b** . . . **330n**, the metadata management system **302**, and other components in the information management system **300** can be connected to one another via a communication network **310**. The communication network **310** can include one or more networks or other connection types including any of following, without limitation: the Internet, a wide area network (WAN), a local area network (LAN), a Storage Area Network (SAN), a Fibre Channel connection, a Small Computer System Interface (SCSI) connection, a virtual private network (VPN), a token ring or TCP/IP based network, an intranet network, a point-to-point link, a cellular network, a wireless data transmission system, a two-way cable system, an interactive kiosk network, a satellite network, a broadband network, a baseband network, a neural network, a mesh network, an ad hoc network, other appropriate wired, wireless, or partially wired/wireless computer or telecommunications networks, combinations of the same or the like. The communication network **310** in some cases may also include application programming interfaces (APIs) including, e.g., cloud service provider APIs, virtual machine management APIs, and hosted service provider APIs.

In an embodiment, the metadata management system **302** is configured to collect metadata corresponding to contents of the storage systems **320a**, **320b** . . . **320n**, and generate and store information relating to characteristics of the stored data and/or metadata. The metadata management system **302** can be present to, for example, store, organize, protect, manage, manipulate, move, analyze, and/or process metadata of data storage systems **320a**, **320b** . . . **320n**. Specifically, the metadata management system **302** may also be configured to generate and store other types of information that generally provides insights into the contents of the storage systems **320a**, **320b** . . . **320n**. The metadata management system **302** can provide a number of benefits including improved storage operations, faster data operation performances, enhanced scalability, or the like. As one specific example which will be discussed below in further detail, the metadata management system **302** can act as a

cache for storing metadata, for analyzing metadata, adding metadata tags, updating metadata tags, or the like.

In certain embodiments, the metadata management system **302** includes a metadata collection system **351** in communication with a metadata store **352** and a classifier **353**.

Generally speaking, the metadata management system **302** may be implemented as a storage system (some combination of hardware and software) that manages, coordinates, and facilitates the transmission of metadata between a client computing device and one or more data storage systems, and/or between the one or more storage systems such that operations related to the metadata management system **302** do not significantly impact performance of other components in the information management system **300**. Moreover, as will be described further, the metadata management system **302** may be configured to make calls to data storage system **320a**, **320b** . . . **320n** and/or receive information from the data storage system **320a**, **320b** . . . **320n** through data network **310**. For example, metadata management system **302** may provide API calls, commands, or other services allowing for the storage, management, and retrieval of metadata, system data blocks, or items. In one embodiment, metadata management system **302** may include or be associated with one or more storage devices, providers, or solutions for the storage of items, system data blocks, or other data.

In an embodiment, the metadata collection system **351** may collect the metadata from data storage systems **320a**, **320b** . . . **320n** and store it in the metadata store **352**. The metadata collected by the metadata collection system **351** may be system metadata, event metadata, scan metadata, or any other type of metadata. System metadata includes metadata collected and stored by the data storage systems **320a**, **320b** . . . **320n** internally using any now or hereafter known methods. Event metadata includes metadata corresponding to an event (or data operation) executed on the data storage systems **320a**, **320b** . . . **320n** and may include, without limitation, information about the data set relating to the event (e.g., file name, location, author, size, or the like); information about the event (e.g., event type, function performed, resulting changes to the data, time of event, or the like); information about the application or client device that performed and/or initiated the event; information about the data storage system on which the event was executed, and/or the like. Scan metadata includes metadata collected by the metadata collection system **351** by externally scanning the contents (e.g., documents, files, objects, images, etc.) of the data storage systems **320a**, **320b** . . . **320n**. System metadata and scan metadata may include, without limitation content metadata that provides information on data objects stored in data storage systems **320a**, **320b** . . . **320n**; volume metadata that provides information on volumes configured in the data storage systems **320a**, **320b** . . . **320n** in which the content is stored; device class metadata that defines the type of storage hardware used for a particular storage pool (the device class may indicate the storage device type and specifies a device type and media management information, such as recording format, estimated capacity, and labeling prefixes); library metadata that provides a further level of abstraction representing a storage entity that contains a media changer in addition to drives and tapes for storing data; and/or the like.

In certain embodiments, each of the data storage systems **320a**, **320b** . . . **320n** may collect and store the system metadata corresponding to contents of the respective data storage systems internally using any now or hereafter known

methods. For example, the data storage systems **320a**, **320b** . . . **320n** may collect metadata when the contents are created, modified, and/or periodically using any now or hereafter known methods. The data storage systems **320a**, **320b** . . . **320n** may transmit the collected system metadata to the metadata collection system **351** (via, for example, an API). The internally collected system metadata may be temporarily and/or permanently stored on the data storage systems **320a**, **320b** . . . **320n**. For example, FIG. 3 illustrates system metadata **321a**, **321b** . . . **321n** stored in data storage systems **320a**, **320b** . . . **320n**, respectively.

Alternatively and/or additionally, the metadata collection system **351** may collect the metadata by performing a periodic scan of one or more of the data storage systems **320a**, **320b** . . . **320n** (“scan metadata”). The metadata collection system may utilize any now or hereafter known techniques to collect the metadata from the data storage systems **320a**, **320b** . . . **320n** (described below). For example, one approach to gathering stored data metadata is by scanning a data storage system from outside using standard client access network protocols such as NFS and CIFS protocols in a NAS file storage context and SCSI in a block storage context. Although not limited to a specific format, the aforementioned metadata can be of a data format referred to as inode. Alternatively, they may be in a data format referred to as NTFS in Windows® OSs. Metadata used in MacOS® may also be used. In certain embodiments, the metadata collection system **330** may use deep learning, machine learning, and/or other methods to parse the contents of the data storage systems **320a**, **320b** . . . **320n** and collect the metadata (e.g., using the IBM Watson™ QA system available from International Business Machines Corporation, or other natural language processing and/or deep learning systems).

Alternatively and/or additionally, the metadata collection system **351** may collect event metadata. In certain embodiments, the data storage systems **320a**, **320b** . . . **320n** may forward event metadata to the metadata collection system **351** upon occurrence of one or more events (e.g., a copy operation, backup operation, an encryption operation, or the like). For example, the metadata collection system **351** may install an event monitoring agent on the data storage systems **320a**, **320b** . . . **320n**. In one or more embodiments, the metadata collection system **351** may configure the data storage systems **320a**, **320b** . . . **320n** to send an event notification along with event metadata every time an event occurs for data residing on a data storage system. Alternatively, the metadata collection system **351** may configure the data storage systems **320a**, **320b** . . . **320n** to send a collection of event notification and/or event metadata periodically.

The metadata collection system **351** may receive streams of log data (e.g., data storage system logs, client device logs) from many sources, convert log entries from the log data into events, and store the events in metadata store **352** based on fields specified in source type definitions (also referred to herein simply as source types). Each event represents a particular log entry. The events that are stored in the metadata store **352** may be based on log entries from various sources and may have different formats. Examples of log entries include simple network management protocol (SNMP) logs, reports from devices and/or applications running on devices, application programming interface (API) call records, information exchange protocols, remote authentication dial-in user service (RADIUS) logs, lightweight directory access protocol (LDAP) logs, security assertion markup language (SAML) messages, and so forth.

These diverse events may all be stored and indexed in the metadata store **352**, which may be a non-homogenous database, in a manner that enables the events to be searched and linked together.

For example, an event monitor agent may include a filter driver program and may be deployed on an input/output port of the data storage system, a read/write port, or data stack and operate in conjunction with a file management program to record events executed on a data storage system. Such operation may involve creating a data structure such as a record or journal of each event. The records may be stored in a journal data structure and may chronicle events in any form or structure (e.g., on an interaction by interaction basis). The journal data structure may include information regarding the type of event that has executed along with certain configurable relevant properties of the data involved in the event. One example of such a monitor program may include Microsoft's Change Journal. Each data storage system may then transmit the event log periodically and/or every time an event occurs to the data storage system. Alternatively and/or additionally, the metadata collection system **351** may periodically consult the recorded interactions for new entries. If new entries exist, the metadata collection system **351** may examine the entries, and if deemed relevant, the entries may be analyzed, parsed, and written to the metadata store **352** as an update.

In some other embodiments, the metadata collection system **351** may also monitor data interactions between the data storage systems **320a**, **320b** . . . **320n** and/or between the client devices **330a**, **330b** . . . **330n** and the data storage systems **320a**, **320b** . . . **320n**, using any suitable monitoring methods, to collect event metadata. For example, the metadata collection system **351** may monitor data interactions by monitoring file system managers associated with each of the data storage systems **320a**, **320b** . . . **320n** (e.g., operating system programs, a FAT, an NTFS, or the like that may be used to manage data movement to and/or from a mass storage device). In another example, the metadata collection system **351** may monitor data interactions by monitoring the network traffic on the communication network **310** using any now or hereafter known methods. In yet another example, the metadata collection system **351** may collect event metadata by interfacing and/or communicating with a virtual file system (VFS) layer that transfers data operation requests between the client devices **330a**, **330b** . . . **330n** and the data storage systems **320a**, **320b** . . . **320n**.

In certain embodiments, the metadata collection system **351** may collect system metadata, scan metadata, and/or event metadata in a manner that duplication of the collected metadata may be minimized. For example, the metadata collection system **351** may analyze the system metadata and may only collect scan metadata for metadata that is not included in the scan metadata and/or to update the system metadata periodically. Similarly, upon initialization the metadata collection system **351** may first collect system metadata and/or scan metadata from the data storage systems **320a**, **320b** . . . **320n** (using one or more methods described above) before starting collection of event metadata. This may be done in order to obtain an accurate picture of the data being scanned and/or to maintain referential integrity within the system.

Duplication may also be prevented by saving only the latest metadata corresponding to a data set. For example, if an event corresponding to a previously scanned data set is registered, the metadata collection system may overwrite the previously stored scan metadata for that data set with the new event metadata (or vice versa).

Events according to certain embodiments are generally data operations executed on the one or more data storage systems such as, without limitation, data migration operations (e.g., copy, backup, archive, email etc.), writing new data on the data storage system, reading data from the data storage system, deletion of data, changing one or more properties of data and/or associated metadata (e.g., rename, access permissions, security, encryption, or the like), printing, or other types of data operations. Such data operations lead to a modification in existing metadata of the corresponding data set and/or creation of new metadata (e.g., when new data set is created). Event metadata, therefore, may also include information relating to changes in the metadata corresponding to a data set.

For example, operations may be available to interact with stored data, including open, write new file (e.g., PUT), write (append to data set), write (modify an existing data set), close, read (e.g., GET), SAVE, RENAME, DELETE, or the like. A PUT operation writes a new object to a storage device of a data storage system or creates an updated version of an existing object on a storage device, and in the latter instance, the previous version may or may not be removed from the storage device. Typically, however, when an updated version of an existing object is written to the memory device, the newer version is identified in metadata as an update (e.g., "version 2"), while older versions (e.g., "version 1") remain stored on the storage device. A DELETE operation is typically associated with writing a new version of an object to a storage device (e.g., via a PUT operation) and indicating a deletion of the old version. Where the old version of the object is physically removed from a storage device, the removal may be hard (e.g., the old version of the object is immediately rewritten as zero-byte version) or soft (e.g., the old version of the object is marked deleted in metadata and later rewritten). In one example of removal, the old version of the object may be cleaned up by an out-of-band process. A GET operation retrieves a representation of an object already stored on a storage device, for instance, in order to read the object.

As discussed above, events may also include data migration operations that involve the copying or migration of data between different locations in the information management system **300** in an original/native and/or one or more different formats. For example, events can include operations in which stored data is copied, migrated, or otherwise transferred from one or more first storage systems to one or more second storage systems, from one or more first storage systems to one or more client devices, and/or within a storage system. Such operations can include by way of example, backup operations, archive operations, information lifecycle management operations such as hierarchical storage management operations, replication operations (e.g., continuous data replication operations), snapshot operations, deduplication or single-instancing operations, auxiliary copy operations, and the like. As will be discussed, some of these operations involve the copying, migration or other movement of data, without actually creating multiple, distinct copies of metadata in the data storage systems itself. Nonetheless, some or all of these operations are referred to as "copy" operations for simplicity.

Backup Operations: A backup operation creates a copy of a version of data (e.g., one or more files or other data units) in a data storage system at a particular point in time. Each subsequent backup copy may be maintained independently of the first. Further, a backup copy in some embodiments is generally stored in a form that is different than the native format, e.g., a backup format. This can be in contrast to the

version in the corresponding data storage system from which the backup copy is derived, and which may instead be stored in a native format of the source application(s). In various cases, backup copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format. For example, a backup copy may be stored in a backup format that facilitates compression and/or efficient long-term storage.

Backup copies can have relatively long retention periods as compared to corresponding data (“primary data”), and may be stored on media with slower retrieval times than primary data and certain other types of secondary copies. On the other hand, backups may have relatively shorter retention periods than some other types of secondary copies such as archive copies (described below). Backups may sometimes be stored at an offsite location. Backup operations can include full, synthetic or incremental backups. A full backup in some embodiments is generally a complete image of the data to be protected. However, because full backup copies can consume a relatively large amount of storage, it can be useful to use a full backup copy as a baseline and only store changes relative to the full backup copy for subsequent backup copies.

For instance, a differential backup operation (or cumulative incremental backup operation) tracks and stores changes that have occurred since the last full backup. Differential backups can grow quickly in size, but can provide relatively efficient restore times because a restore can be completed in some cases using only the full backup copy and the latest differential copy.

An incremental backup operation generally tracks and stores changes since the most recent backup copy of any type, which can greatly reduce storage utilization. In some cases, however, restore times can be relatively long in comparison to full or differential backups because completing a restore operation may involve accessing a full backup in addition to multiple incremental backups.

Any of the above types of backup operations can be at the volume-level, file-level, or block-level. Volume level backup operations generally involve the copying of a data volume (e.g., a logical disk or partition) as a whole. In a file-level backup, the information management system 300 may generally track changes to individual files at the file-level, and includes copies of files in the backup copy. In the case of a block-level backup, files are broken into constituent blocks, and changes are tracked at the block-level. Upon restore, the information management system 300 reassembles the blocks into files in a transparent fashion.

Far less data may actually be transferred and copied to secondary storage devices during a file-level copy than a volume-level copy. Likewise, a block-level copy may involve the transfer of less data than a file-level copy, resulting in faster execution times. However, restoring a relatively higher-granularity copy can result in longer restore times. For instance, when restoring a block-level copy, the process of locating constituent blocks can sometimes result in longer restore times as compared to file-level backups. Similar to backup operations, the other types of secondary copy operations described herein can also be implemented at either the volume-level, file-level, or block-level.

Archive Operations: Because backup operations generally involve maintaining a version of the copied data in primary data and also maintaining backup copies in secondary storage device(s), they can consume significant storage capacity. To help reduce storage consumption, an archive operation

according to certain embodiments creates a secondary copy by both copying and removing source data. Or, seen another way, archive operations can involve moving some or all of the source data to the archive destination. Thus, data satisfying criteria for removal (e.g., data of a threshold age or size) from the source copy may be removed from source storage. Archive copies are sometimes stored in an archive format or other non-native application format. The source data may be primary data or a secondary copy, depending on the situation. As with backup copies, archive copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format. In addition, archive copies may be retained for relatively long periods of time (e.g., years) and, in some cases, are never deleted. Archive copies are generally retained for longer periods of time than backup copies, for example. In certain embodiments, archive copies may be made and kept for extended periods in order to meet compliance regulations.

Moreover, when primary data is archived, in some cases the archived primary data or a portion thereof is deleted when creating the archive copy. Thus, archiving can serve the purpose of freeing up space in the primary storage device(s). Similarly, when a secondary copy is archived, the secondary copy may be deleted, and an archive copy can therefore serve the purpose of freeing up space in secondary storage device(s). In contrast, source copies often remain intact when creating backup copies.

Snapshot Operations: Snapshot operations can provide a relatively lightweight, efficient mechanism for protecting data. From an end-user viewpoint, a snapshot may be thought of as an “instant” image of the primary data at a given point in time, and may include state and/or status information relative to an application that creates/manages the data. In one embodiment, a snapshot may generally capture the directory structure of an object in primary data such as a file or volume or other data set at a particular moment in time and may also preserve file attributes and contents. A snapshot in some cases is created relatively quickly, e.g., substantially instantly, using a minimum amount of file space, but may still function as a conventional file system backup.

A “hardware snapshot” (or “hardware-based snapshot”) operation can be a snapshot operation where a target storage device (e.g., a primary storage device or a secondary storage device) performs the snapshot operation in a self-contained fashion, substantially independently, using hardware, firmware and/or software residing on the storage device itself. For instance, the storage device may be capable of performing snapshot operations upon request, generally without intervention or oversight from any of the other components in the information management system 300. In this manner, hardware snapshots can off-load other components of information management system 300 from processing involved in snapshot creation and management.

A “software snapshot” (or “software-based snapshot”) operation, on the other hand, can be a snapshot operation in which one or more other components in information management system 300 implement a software layer that manages the snapshot operation via interaction with the target storage device. For instance, the component implementing the snapshot management software layer may derive a set of pointers and/or data that represents the snapshot. The snapshot management software layer may then transmit the same to the target storage device, along with appropriate instructions for writing the snapshot.

Some types of snapshots do not actually create another physical copy of all the data as it existed at the particular point in time, but may simply create pointers that are able to map files and directories to specific memory locations (e.g., to specific disk blocks) where the data resides, as it existed at the particular point in time. For example, a snapshot copy may include a set of pointers derived from the file system or an application. In some other cases, the snapshot may be created at the block-level, such as where creation of the snapshot occurs without awareness of the file system. Each pointer points to a respective stored data block, so that collectively, the set of pointers reflect the storage location and state of the data block (e.g., file(s) or volume(s) or data set(s)) at a particular point in time when the snapshot copy was created.

Once a snapshot has been taken, subsequent changes to the file system typically do not overwrite the blocks in use at the time of the snapshot. Therefore, the initial snapshot may use only a small amount of disk space needed to record a mapping or other data structure representing or otherwise tracking the blocks that correspond to the current state of the file system. Additional disk space is usually required only when files and directories are actually later modified. Furthermore, when files are modified, typically only the pointers which map to blocks are copied, not the blocks themselves. In some embodiments, for example in the case of “copy-on-write” snapshots, when a block changes in primary storage, the block is copied to secondary storage or cached in primary storage before the block is overwritten in primary storage, and the pointer to that block changed to reflect the new location of that block. The snapshot mapping of file system data may also be updated to reflect the changed block(s) at that particular point in time. In some other cases, a snapshot includes a full physical copy of all or substantially all of the data represented by the snapshot.

A snapshot copy in many cases can be made quickly and without significantly impacting primary computing resources because large amounts of data are not copied or moved. In some embodiments, a snapshot may exist as a virtual file system, parallel to the actual file system. Users in some cases gain read-only access to the record of files and directories of the snapshot. By electing to restore primary data from a snapshot taken at a given point in time, users may also return the current file system to the state of the file system that existed when the snapshot was taken.

**Replication Operations:** Another type of secondary copy operation is a replication operation. Some types of secondary copies are used to periodically capture images of primary data at particular points in time (e.g., backups, archives, and snapshots). However, it can also be useful for recovery purposes to protect primary data in a more continuous fashion, by replicating the primary data substantially as changes occur. In some cases a replication copy can be a mirror copy, for instance, where changes made to primary data are mirrored or substantially immediately copied to another location (e.g., to secondary storage device(s)). By copying each write operation to the replication copy, two storage systems are kept synchronized or substantially synchronized so that they are virtually identical at approximately the same time. Where entire disk volumes are mirrored, however, mirroring can require significant amount of storage space and utilizes a large amount of processing resources.

**Deduplication/Single-Instancing Operations:** Another type of data movement operation is deduplication or single-instance storage, which is useful to reduce the amount of data within the system. For instance, some or all of the

above-described secondary storage operations can involve deduplication in some fashion. New data is read, broken down into portions (e.g., sub-file level blocks, files, etc.) of a selected granularity, compared with blocks that are already stored, and only the new blocks are stored. Blocks that already exist are represented as pointers to the already stored data.

**Information Lifecycle Management and Hierarchical Storage Management Operations:** In some embodiments, files and other data over their lifetime move from more expensive, quick access storage to less expensive, slower access storage. Operations associated with moving data through various tiers of storage (e.g., as shown in FIG. 3) are sometimes referred to as information lifecycle management (ILM) operations.

One type of ILM operation is a hierarchical storage management (HSM) operation. A HSM operation is generally an operation for automatically moving data between classes of storage devices, such as between high-cost and low-cost storage devices. For instance, an HSM operation may involve movement of data from primary storage devices to secondary storage devices, or between tiers of the same storage devices. With each tier, the storage devices may be progressively relatively cheaper, have relatively slower access/restore times, etc. For example, movement of data between tiers may occur as data becomes less important over time.

In some embodiments, an HSM operation is similar to an archive operation in that creating an HSM copy may (though not always) involve deleting some of the source data, e.g., according to one or more criteria related to the source data. For example, an HSM copy may include data from primary data or a secondary copy that is larger than a given size threshold or older than a given age threshold and that is stored in a backup format. Often, and unlike some types of archive copies, HSM data that is removed or aged from the source copy is replaced by a logical reference pointer or stub. The stub may also include some metadata associated with the corresponding data, so that a file system and/or application can provide some information about the data block and/or a limited-functionality version (e.g., a preview) of the data block. According to one example, files are generally moved between higher and lower cost storage depending on how often the files are accessed. An HSM copy may be stored in a format other than the native application format (e.g., where the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format). In some cases, copies which involve the removal of data from source storage and the maintenance of stub or other logical reference information on source storage may be referred to generally as “on-line archive copies”. On the other hand, copies which involve the removal of data from source storage without the maintenance of stub or other logical reference information on source storage may be referred to as “off-line archive copies”.

It will be understood to those skilled in the art that it is possible to employ “event” definitions that may capture a relatively broad or narrow set of data operations executed on a data storage system, allowing a user to customize the metadata collection system 351 to meet certain metadata collection goals. Such “event” definitions may define or describe data movement, changes, manipulations or other operations or interactions that may be of interest to a system user or administrator (e.g., any operation that “touches” data may be recorded along with the action or operation that caused the interaction (e.g. read, write, copy, parse, or the



like). Moreover, change definitions may evolve over time or may be dynamic based on the entries sent to the metadata store **352**. For example, if expected results are not obtained, change definitions may be modified or additional definitions used until appropriate or desired results are obtained. This may be accomplished, for example by globally linking certain libraries of “event” definitions and selectively enabling libraries on a rolling basis until acceptable results are achieved. This process may be performed after the initial activation of the metadata collection system **351** and periodically thereafter, depending on changing needs or objectives.

Moreover, in some embodiments, the system may support the use of “user data tags” that allow certain types of information stored in the data storage systems **320a**, **320b** . . . **320n** to be tagged so they may be identified and tracked throughout the system. As such, if a data block that includes a user data tag is touched, an event log is recorded and/or sent or collected by the metadata collection system **351**. For example, a user may designate a particular type of data or information such as project information, or information shared between or accessed by particular group of users to be tracked across the system or through various levels of storage. This may be accomplished through a user interface that allows a user to define certain information to be tagged, for example, by using any available attribute within the system such as those specified above with respect to the classification agent or filter used in the system. In some embodiments, the user may define one or more tags using these or other attributes which may be further refined by combining them through the use of logical or Boolean operators to define a certain tag expression.

For example, a user may define a certain user data tag by specifying one or more criteria to be satisfied such as certain system users, a certain data permission level, a certain project, combinations of the same or the like. These criteria may be defined using logical operators such as AND or OR operators to conditionally combine various attributes to create a condition that defines a tag. In certain embodiments, information satisfying that criteria may be tagged and tracked within the system. For example, the metadata store **352** may contain entries keeping track of entries satisfying the tag criteria along with information relating to the types of operations performed on the information as well as certain metadata relating to the data content and its location in the data storage systems **320a**, **320b** . . . **320n**. This allows the system to search the metadata store **352** at a particular level of storage for the information, and quickly locate it within mass storage device for potential retrieval.

Referring back to FIG. 3, metadata (system metadata and/or event metadata) collected by the metadata collection system **351** is stored in the metadata store **352** outside of the data storage systems **320a**, **320b** . . . **320n**. In one or more embodiments, the metadata store **352** may be any type of data structure that allows for easy and efficient searching of the stored metadata. Examples may include, without limitation, relational database storage (e.g., SQL databases), key-value type storages (e.g., noSQL databases), columnar storages (e.g., parquet), or the like. The metadata store **352** may also include an index associated with each piece of metadata and stored with the metadata. The index may contain information such as each of the locations where the data set corresponding to the metadata is located, user access information describing which users are permitted to view the contents of the data set, type of data structure corresponding to the data set, or the like. The content index may be used

to facilitate search and retrieval of a data set corresponding to metadata, such as in response to a user request to restore a particular file.

In one or more embodiments, the metadata store **352** may be a data structure in the form of a NoSQL (“Not-only-Structured-Query-Language”) database. In one embodiment, the metadata store **352** is implemented using a NoSQL database that uses a key-value store, a document store, and/or a wide column store. Specifically, event metadata collected by the metadata collection system **351** may be stored in a NoSQL type database structure. A NoSQL database may be, by way of example, Cloudant® Apache Cassandra™, Object Storage, Apache HBase™, Hazelcast®, etc.

A NoSQL database provides a mechanism for storage and retrieval of data that is modeled in means other than the tabular relations used in relational databases. Typical motivations for this approach include simplicity of design, horizontal scaling, and finer control over availability. NoSQL databases have features of self-organizing, self-managing, low cost, high scalability, high concurrency, simply query relation, and so on. To compare a NoSQL database to a relational database, a form in the relational database usually stores a formatted data structure, and components of all entry fields are the same. Even if not every entry needs all fields, the relational database will allocate all fields to each entry. Such structure can potentially cause a performance bottleneck in a relational database. On the other hand, a NoSQL database typically carries out storage with a Key/Value pair, and its structure is not fixed. Each entry can have different fields, and each entry can add some key value pairs of its own according to a requirement so as not to be limited to the fixed structure and to thereby reduce some time and space overheads.

In an embodiment, an “entry” corresponding to an event may be a record in a NoSQL database, and can also be regarded as a data object instance in the NoSQL database. Each entry can possess a unique identifier (ID), and can comprise zero or more Key/Value pairs. Usage examples include storing millions of data records as key-value pairs in one or a few associative arrays. A key-value pair is a fundamental data representation in computing systems and applications, in which all or part of the data model may be expressed as a collection of tuples <attribute name, value>, for which each element is a key-value pair. An associative array is an unordered list of unique attributes with associated values. Such organization is particularly useful for statistical or real-time analysis of growing lists of data elements. According to an embodiment of the present invention, a pre-defined specifier can be used to distinguish between individual Key/Value pairs. For example, different Key/Value pairs are distinguished by a comma. Meanwhile, the “key” and the “value” within each Key/Value pair can be separated by a pre-defined delimiter, for example, a colon, thus the key in a Key/Value pair can be determined from the Key/Value pair according to the delimiter. At the same time, the “value” in a Key/Value pair can be extended by a pre-defined extension symbol, for example, square brackets which can be used to represent that the “value” in a Key/Value pair comprises more than two attributes. Each attribute in the more than two attributes can either be a real “value”, or be a Key/Value pair in which the “value” can continue to comprise one or more attribute.

For example, metadata (event metadata, scan metadata and system metadata) may include one or more of the following keys and corresponding values:

Keys	Values
Data Storage System	Identity of the Data Storage System on which the event was executed
Time Stamp	Time of event
File	File Name of the data set (or “file”) or which the event was executed
Title	Title of the data set
Type	Type of content (e.g., document (WORD, PDF, etc.)/ image (JPEG, GIFF, etc./video/email/web link/blog/or the like)
Size	Size of file/size of data corresponding to event execution
Time of creation	File Creation time
Owner	Creator of file (user name, client device, affiliation, department, or other user information etc.)
Author	Event executor (user name, client device, affiliation, or other user information etc.)
Event type	Copy/Read/Write/Modify/Delete/Print/Email/etc.
Path	Data path corresponding to event (e.g., copy {F1} to {F2})
Text	Text Excerpts
Facets	Facets (e.g., social security number, patient name, etc.)
Tags	Keywords (e.g., financial, personal, sensitive, medical, etc.) and/or previously associated tags

It should be noted that the above key values are provided as way of example only and should not be considered limiting.

It will be understood to those skilled in the art that while other types of storing formats are not described here in detail, the metadata may be stored in other formats as well (e.g., SQL, Parquet, etc.).

In certain embodiments, the metadata management system 302 also includes a classifier 353 configured to apply custom tags (interchangeably referred to as tags) to the metadata before and/or after insertion by the metadata collection system 351 into the metadata store 352. In certain embodiments, the facets (described below) associated with the metadata that are indicative of the actual content in the data storage systems 320a, 320b . . . 320n that corresponds to the tagged metadata and/or the metadata itself may be used as tags.

In some embodiments, the classifier 353 analyzes characteristics, content, format, etc. of the metadata (and not the data itself) to add tags to the metadata. This provides enhanced search and management capabilities for data discovery and other purposes. The custom tags, in certain embodiments, significantly reduce the amount of time required to obtain information by reducing and/or substantially eliminating the need to obtain information directly from the source data. The tags can be used to identify files or other data blocks in the data storage systems 320a, 320b . . . 320n having pre-defined content (e.g., user-defined keywords or phrases, other keywords/phrases that are not defined by a user, etc.), and/or metadata (e.g., email metadata such as “to”, “from”, “cc”, “bcc”, attachment name, received time, etc.).

For example, assume a system administrator desires to identify data sets that a certain user has interacted with, that contain content including certain keywords, content having characteristics, etc. Rather than search each file in each directory and/or all the metadata content, which can be a very time consuming process (especially when the data blocks reside on multiple storage devices or the metadata store includes large volumes of data in a schema-less database format), the system administrator may search the custom tags in the metadata store 352 to identify metadata that is associated with tags corresponding to the user, keywords and/or characteristics (by for example, defining a query), and may then look up data sets associated with that metadata.

Moreover, in certain embodiments, use of the custom tags in a metadata store 352 where the custom tags do not reside on a data storage system itself for satisfying data searches or queries may also reduce the involvement of network resources in this process, substantially reducing the processing burden on the host system. For example, as described above, if an administrator desires to identify certain data sets, querying the metadata store 352 rather than the file system virtually removes the host system from the query process (e.g., no brute force scanning of directories and files in the data storage systems is required), allowing the host system to continue performing host tasks rather than be occupied with search tasks.

The classifier 353 may apply the tags by analyzing the collected metadata (before and/or after insertion into the metadata store 352) and applying one or more user defined policies. Alternatively and/or additionally the classifier 353 may apply the tags automatically by analyzing the collected metadata (before and/or after insertion into the metadata store 352), and applying one or more classification rules automatically derived by the classifier 353 (e.g., based on machine learning, deep learning, text annotators, or the like). In yet another embodiment, the classifier 353 may analyze the content in the data storage systems 320a, 320b . . . 320n, and apply one or more classification rules automatically derived by the classifier 353 (e.g., based on machine learning, deep learning, text annotators, or the like) to tag collected metadata associated with that content. Different methods for applying the custom tags are described below in detail.

In an embodiment, where the classifier 353 may apply the tags by analyzing the metadata and applying one or more user defined policies, the classifier 353 may include and/or may be in communication with a policy engine 355. The policy engine 355 may include a set of user configurable policies. In these examples, a policy is a set of classification rules. The policy may also include any other data or parameters in addition to the set of classification rules that can be used to interpret how the metadata tags are to be assigned.

For example, a policy may include multiple rules for classifying metadata and/or may specify the tag(s) to be applied if the rule conditions are satisfied. For example, a user may specify a tag to be associated with a metadata based on one or more characteristics of the metadata.

In certain embodiments, the set of rules for a tag can be defined in a tag definition that is typed directly into a user interface program (e.g., a REST API, SDK, or the like) and written into the policy engine. In an alternative embodiment, the tag definition can be represented in a definition file. If a definition file is used, it can use the XML markup language or any document structure. In an embodiment, a user may create new customized tags. Alternatively and/or additionally, the user may be presented a list of pre-defined tags and the user may choose and/or modify such tags. The tags may be relevant to a user, task, business objective, or the like.

A user may define a classification policy by indicating criteria, parameters or descriptors of the policy via a graphical user interface that provides facilities to present information and receive input data, such as a form or page with fields to be filled in, pull-down menus or entries allowing one or more of several options to be selected, buttons, sliders, hypertext links or other known user interface tools for receiving user input. For example, a user may define tags “confidential” and “access level 2” if the metadata includes certain keywords (e.g., “confidential,” or “privileged”) and/or are associated particular flags (e.g., in metadata identifying a document or email as personal, confidential, etc.).

A policy defines a particular combination of rules, such as users who have created, accessed or modified a document or data block; file or application types; content or metadata keywords; clients or storage locations; dates of data creation and/or access; review status or other status within a workflow (e.g., reviewed or un-reviewed); modification times or types of modifications; and/or any other data attributes. A policy may also be defined using tags already associated with the metadata. For example, a rule may classify all metadata associated with an already assigned tag (e.g., “project”), and apply a second tag (e.g., a second tag “inactive” that corresponds to the status of projectX).

The various rules used to define a policy may be combined in any suitable fashion, for example, via Boolean operators, to define a complex policy. As an example, an E-discovery policy might define a tag “privileged” that is associated with documents or data blocks that (1) were created or modified by legal department staff (i.e., owner or creator in the metadata is associated with a defined name), (2) were sent to or received from outside counsel via email, and/or (3) contain one of the following keywords: “privileged” or “attorney,” “counsel”, or other terms.

Another type of tag which may be added is an entity tag. An entity tag may be, for example, any content that matches a defined data mask format. Examples of entity tags might include, e.g., social security numbers (e.g., based on a rule that any numerical content matching the formatting mask XXX-XX-XXXX), credit card numbers (e.g., based on a rule content having a 13-16 digit string of numbers), SKU numbers, product numbers, etc.

Policies may, in certain embodiments, include one or more of the following classification rules for assigning tags to the metadata:

- i. frequency with which metadata and/or corresponding has been or is predicted to be used, accessed, or modified;
- ii. size of metadata;
- iii. user information that created, accessed, modified, or otherwise utilized content corresponding to the metadata (e.g., owner, creator, author, etc.): based on user name, user affiliation, user access level, etc.
- iv. time-related factors (e.g., aging information such as time since the creation or modification of a metadata);
- v. the identity of applications, client devices and/or other computing devices that created, accessed, modified, or otherwise utilized content corresponding to the metadata;
- vi. a relative sensitivity (e.g., confidentiality) of a data block, e.g., as determined by its content and/or metadata;
- vii. the current or historical storage capacity of various storage devices;
- viii. the current or historical network capacity of network pathways connecting various components within the storage operation cell;
- ix. access control lists or other security information;
- x. already existing tags associated with the metadata; and/or the content of metadata (e.g., keywords, tags, etc.).

In an embodiment, the policies defined by the user may also be in a key-value form (if metadata is stored in a noSQL database format), and the classifier 353 may search for the key-value in the metadata for application of a tag defined by the policy. In certain embodiments, the user may be prompted to define keys that are used by the metadata store 352 for sorting and storing the metadata. Optionally, a user may define any key, and the classifier 353 may apply the tag

if a metadata entry includes either an exact match for the key-value included in the policy and/or if a key-value is similar to the key-value included in the policy. For example, if the policy includes a key-value “owner:bob” in the policy, the classifier may apply the tag if a metadata entry includes either an exact match for the key-value (i.e., owner:bob) included in the policy and/or a similar key-value (e.g., creator:bob; author:bob; etc.). In one or more embodiments, tagging guidance may be provided to a user for which the accuracy may increase with use, as the history of the user select tags builds up (e.g., using a feedback loop that provides user with guidance relating to most common keys in the metadata store 352).

For example, consider a user is creating documents for different enterprises having names “/usr/documents/IBM/”; “/usr/documents/Company2/”; “/usr/documents/Company3/” etc. in a data storage system. The metadata collection system will receive event metadata corresponding to the creation of the documents that will include the document names, formats, sizes, date of creation, etc. The classifier 353 may add company name tags to the event metadata based on a user defined policy including rules for classifying metadata based on company name (by extracting the company name from the file names included in the event metadata). Another tag describing the document format may also be added (e.g., docx, pdf, etc.).

In certain embodiments, tagging based on user-defined policies may be performed for event metadata upon receipt of event metadata for each event by the metadata collection system 351 (e.g., on a first-in, first-out flow) before insertion into the metadata store 352. Alternatively and/or additionally, tagging based on user-defined policies may be performed for event metadata periodically; or upon occurrence of certain conditions (e.g., receipt of user instructions) before (e.g., in a queue) and/or after insertion into the metadata store 352. Similarly, tagging based on user-defined policies may be performed for scan metadata upon receipt of scan metadata every time a scan is performed by the metadata collection system 351 (e.g., on a first-in, first-out flow) before insertion into the metadata store 352. Alternatively and/or additionally, tagging based on user-defined policies may be performed for scan metadata periodically; or upon occurrence of certain conditions (e.g., receipt of user instructions) before (e.g., in a queue) and/or after insertion into the metadata store 352.

In certain embodiments, the classifier 353 may apply the tags automatically by analyzing the metadata collected by the metadata collection engine 351 before and/or after insertion into the metadata store 352, and applying one or more classification rules automatically derived by the classifier 353. It should be noted that data included in certain components of metadata contain meaningful information that is indicative of important facets of the contents corresponding to the metadata, and which can be extracted and analyzed without extracting information from the contents itself (e.g., files and objects). For example, file system path information, file name, object bucket/container information, object name, owner information, event information, or the like included in the metadata may include information about important facets of the contents corresponding to the metadata and/or the metadata itself. A facet may comprise a specific type of information about content to be determined from metadata components and may include words, phrases, or other data descriptors identifying unique features of a document/data/content/metadata/etc. For instance, a facet may comprise a characteristic of metadata or content type determined by text analytics (e.g., sensitive, privileged, or

the like). Other examples of such facets may include, without limitation, organization names, content type, location, user information, or the like. Facets may be words, phrases, or other data extracted directly from the metadata (e.g., owner name, entity name, document type, etc.) and/or words, phrases, or other descriptors derived based on information included in the metadata.

In an embodiment, the classifier **353** may include a data miner **356** (e.g., text miner, image miner, audio miner, video miner, and/or the like) to apply data analytics to the components of the metadata to determine facets associated with the content corresponding to the metadata. Data analytics provides techniques to convert textual, audio, video or speech data into structured data by extracting information e.g., person names, addresses, etc. and classifying content into categories based on the data and content (i.e., facets).

The data miner **356** may comply with the Unstructured Information Management Architecture (UIMA), and include such annotators as a language identification annotator to identify the language of metadata; a linguistic analysis annotator to apply linguistic analysis to the metadata; a dictionary lookup annotator to match words and synonyms from a dictionary with words in the content of the metadata and to associate keywords with user-defined facets; a named entity recognition annotator to extract person names, locations, and company names; or the like. It should be understood that any text analytic technologies similar to UIMA may be employed to accomplish the techniques described herein. For example, other off-the-shelf analytics applications or custom software and/or hardware may be used instead of, or in addition to, UIMA.

As is known to those skilled in the art, UIMA developed by IBM Corporation (Armonk, N.Y.) is an open platform for creating, integrating and deploying unstructured information management solutions from combinations of semantic analysis and search components to discover patterns. It allows easy authoring of annotators, such as the expression of the format of telephone numbers, or dates, or meeting rooms. Then, given a set of text documents, the UIMA tool applies the various annotators authored, thereby automatically annotating segments of text by different annotations as authored. IBM product platforms that expose the UIMA interfaces include the OmniFind Enterprise Edition and Analytics Edition. The former features UIMA for building full-text and semantic search indexes, and the latter deploys UIMA for information extraction and text analysis. The annotators may be driven off of entity spotting, using Information Extraction (IE) techniques, and/or using natural language mining (NLM) techniques.

For example, in certain embodiments, the classifier **353** may pre-determine categories such as “sensitive”, “enterprise”, “medical content”, or the like. Further, each category may be associated with a particular annotator in the data miner **356**. An annotator can be any combination of dictionaries, parsing rules, character rules, language identification, semantic analysis, and the like. For example, a “Medical Content” category may include dictionaries for topics such as surgery, patient, physician, medicine, medical, and the like. In certain embodiments, custom dictionaries may also be leveraged to classify data. For example, a custom dictionary that contains terms found in the code development environment may be used and may include terms such as ISO, src, software development project names, acronyms associated with development environments (e.g., PMR, etc.). Annotators may also be defined that derive relationships in between terms as well.

It should be noted that dictionary entries may be in the form of a noun, verb, list of causes, and a causation time frame. The dictionary may also contain other data such as parts of speech (e.g., adjectives, adverbs, etc.), phrases, etc. The dictionary may also contain more complex grammar-like constructs. For example, the dictionary may contain noun alternatives and plurals, verb conjugations, and conjunctions or other Boolean terms (e.g., not, or, and, and exclusive-or). The dictionary may be in any format (e.g., plain text, relational database tables, nested XML code, etc.). Any number of dictionaries may be used for analysis. Hence, the system enables the use of dictionary based annotators, entity based annotators or extraction (e.g., organizations, people, location, etc.), in addition to the ability to create custom annotators based on linguistic nuances and terms of specific industries as part of the classification process.

In certain embodiments, the facets extracted may be used to tag the metadata and/or identify a different tag, which may be used for retrieval of content corresponding to the data. Alternatively and/or additionally, facets may be added to a metadata entry in the metadata store **352** before and/or after insertion into the metadata store **352**.

In one or more embodiments, the classifier **353** may also include a rules engine **357** configured to analyze the facets extracted by the text miner **356** to classify the content corresponding to the metadata, and apply custom tags to the metadata (discussed below). Specifically, a tag may be added based on one or more characteristics of the extracted facets. For example, in the above dictionary examples for code development environment, if International Organization for Standardization (ISO) and problem management report (PMR) words are present (extracted facets), the rules engine **357** may tag the metadata with “sensitive” tag indicating that the corresponding content includes sensitive data. Similarly, if the extracted facets include person names that are associated with an entity, the rules engine **357** may tag the metadata with “entity name” tag indicating that the corresponding content is associated with the entity. As such, the classifier **353** may add ‘tags’ to entries in the metadata store **352** for identifying various matched concepts (based on facet analysis) and/or to map the entries to standardized resources. In another example, if the extracted facets include medical information and a patient name, the rules engine **357** may apply tags such as confidential, personal information, etc.

The rules engine **357** may utilize a process of recognizing the relationships, predicates, or dependencies of components of the metadata and/or facets, and thereby extract new, hidden, indirect, or detailed structural information to classify the extracted metadata and apply a tag. For example, the rules engine **357** may include an NLP component that evaluates the extracted facets and may determine whether the facets include a term from a given dictionary in relationship (e.g., immediately followed by) with a term from a related dictionary. If the metadata component satisfies that parsing rule, then the NLP component determines a likelihood that content is about the subject matter based on correlations. If the likelihood exceeds a threshold, the NLP component applies a tag to the event metadata that is indicative of a property of the subject matter. For example, if the metadata includes a file system path that belongs to an entity recognized as a law firm, and the file name includes “non-infringement opinion”; the classifier may extract facets such as legal document, opinion, client information, etc.; and the rules engine may analyze the facets and their relationships to identify tags such as confidential, attorney

client privileged, work product, legal opinion, or the like. In certain embodiments, the facets may also be used as tags.

While the above disclosure describes using text miners to extract facets from the metadata, other methods such as deep semantic relationship detection and analysis, machine learning, or the like are within the scope of this disclosure.

While the above disclosure describes adding tags to the metadata by analyzing the contents of the metadata itself, the disclosure is not so limiting. In certain embodiments, the system may also analyze specific contents of the data storage systems **320a**, **320b** . . . **320n** to extract facets of the contents, and use the facets to add one or more tags to the metadata corresponding to the content in the metadata store **352**. In an embodiment, the metadata management system **302** may include a facet extraction engine **358** configured to analyze content stored in the data storage systems **320a**, **320b** . . . **320n** to extract facets. Facets may include various dimensions of the content such as, without limitation, users, keywords, time stamps, entity names, document type, or the like. In some embodiments, the extracted facets may be utilized as tags in the in the metadata store **352**. Alternatively and/or additionally, the rules engine **357** may use the extracted facets to identify one or more tags to be added to the metadata (as described above).

In certain embodiments, the facet extraction engine **358** may extract facets by analyzing the contents of files, objects, etc. stored in the data storage systems **320a**, **320b** . . . **320n**. Alternatively and/or additionally, the facet extraction engine **358** may extract facets from, for example, metadata associated with the files, objects, file systems, object containers, storage devices, data storage systems, or the like, that resides in the data storage systems **320a**, **320b** . . . **320n**. For example, the facet extraction engine **358** may extract facets from file headers, object names, etc. that may include metadata (e.g., size, date of creation, date of modification, author, entity, storage location, etc.) about the file or object and its contents. For example, most image file headers store information about image format, size, resolution and color space, and optionally authoring information such as who made the image, when and where it was made, what camera model and photographic settings were used (Exif), and so on. Such metadata may be used by the facet extraction engine **358** to extract facets of the contents of the file or object itself.

In certain embodiments, the contents of the data storage systems **320a**, **320b** . . . **320n** that may be analyzed for extracting content facets in response to a data operation to be performed on the contents of a data storage system. For example, if the information system **300** receives a request to perform a data operation on certain data sets stored in one or more of the data storage systems **320a**, **320b** . . . **320n**, the system may analyze those data sets to extract content facets corresponding to the data sets before performing the data operation. For example, if a user requests that sensitive content included in a storage subsystem of data storage system **320a** must be encrypted, the facet extraction engine **358** may analyze content of that storage subsystem, without limitation, entity information, user information, keywords, or the like that indicate that the content is sensitive. Other data operation requests may include, without limitation, backup, archive, deduplication, or the like.

In one or more embodiments, the facet extraction engine **358** may perform facet extraction using supervised learning, unsupervised learning and/or deep inspection methods (including, for example, the data mining methods described above). For example, the facet extraction engine **358** may utilize the supervised and/or unsupervised learning methods

for named entity extraction and/or classification. Named entity recognition and classification are important aspects of information extraction to identify information units such as people, organizations, location names, and numeric expressions for time, money and numbers from unstructured text. Typically, information units or numeric expressions are first extracted out as named entities from the unstructured text (i.e., named entity recognition), followed by learning a function from an entity to its type, which is selected from predefined categories such as: People, Organizations, Locations, Products, Genes, Compounds, and Technologies, etc. (i.e., named entity classification). There are several kinds of learning methods depending on the availability of training examples. Supervised learning methods infer rules from positive and negative examples of named entities over a large collection of annotated documents for each entity type. Supervised learning requires a large annotated corpus and thus is impractical where manually generated labels are not available or are difficult to generate. Unsupervised learning methods apply clustering technology to automatically gather entities from clusters.

For example, a machine learning system **358** may be IBM Watson™ system that is an application of advanced natural language processing, information retrieval, knowledge representation and reasoning, and machine learning technologies to the field of open domain question answering. The IBM Watson™ system is built on IBM's DeepQA™ technology used for hypothesis generation, massive evidence gathering, analysis, and scoring. DeepQA™ takes an input question, analyzes it, decomposes the question into constituent parts, generates one or more hypothesis based on the decomposed question and results of a primary search of answer sources, performs hypothesis and evidence scoring based on a retrieval of evidence from evidence sources, performs synthesis of the one or more hypothesis, and based on trained models, performs a final merging and ranking to output an answer to the input question along with a confidence measure.

Referring back to FIG. 3, in certain embodiments, the information management system **302** may also include and/or may be in communication with a security module **360** configured to identify the sensitivity level of data stored in the data storage systems **320a**, **320b** . . . **320n** and perform the appropriate data protection actions based on the identified sensitivity level.

In an embodiment, the security module **360** may automate the identification of confidential, or sensitive, data in a data storage system and further automate performance of certain actions based on the identified sensitivity level. At its core, the approach describes a process with the components of a knowledge base (a "Corpora") that is ingested from different structured or unstructured information available about the data storage system. Moreover, the corpora (or corpus) has itself been "ingested" by a set of pre-processing steps that use NLP to analyze the content and transform it in a format adapted to the DeepQA Analysis engines. An artificial intelligence (AI) engine, such as the IBM Watson system, that analyzes the data in light of the knowledge base to identify the sensitivity level of the data. Further, based on the sensitivity level identified, the approach performs an appropriate action.

FIG. 4 is a block diagram illustrating the various components of an example security module **360**. In an embodiment, the security module may include a knowledge base **402**, a policies engine **404**, an inference engine **406**, a data handling agent **408**, and an alert miner **410**.

The knowledge base **402** (also known as a “Corpora”) may include both internal and externally available data sources. Examples of such data source may include, without limitation, confidentiality policies data store, confidential documents data store, code names data store, trade secrets data store, personal identifiable information (PII) data store, financial information data store, personal information data store, medical data store, pattern matching data store (e.g., formats for dates, social security numbers, bank account numbers, or the like). Such data stores may be configurable and may be selected by an organization or enterprise based on its missions, policies, use-cases, etc. For example, a hospital may define and/or configure data stores relating to PII, medical information, financial information, personal information, code names, confidentiality policies, or the like.

In certain embodiments, the policies engine **404** may store one or more rules for performing various security actions and/or preferred methods for performing the security actions on content or data depending upon its sensitivity level. Examples of such security actions and/or preferred methods for performing the security actions may include, without limitation, encryption; redaction of information; backup; archiving; duplication; application of access controls; application of dissemination controls; inclusion and exclusion of partitions, folders, files, file extensions or other divisions of information; specifying encryption protocols and/or levels for encryption; applying and configuring operating system protection rules; applying and configuring user authentication and access recovery mechanisms; applying and configuring device usage controls; auditing of system or user behavior; applying and configuring firewall rules; applying and configuring rules regarding anti-virus, anti-spam and anti-spyware; and applying and configuring rules regarding the provisioning, management and integration with third party applications such as email and synchronization tools; or the like.

The policies engine **404** may automatically define the sensitivity levels and their corresponding security actions based on, for example, business needs, types of sensitive information encountered, potential uses and/or misuses of the sensitive information, available computing resources, or the like. Alternatively, a user (e.g., an enterprise administrator) may provide the different rules for associating the sensitivity levels and the security actions.

In an embodiment, the inference engine **406** may receive data (metadata tags and/or facets) from the metadata store **352** and/or metadata collection system **351** before insertion into the metadata store **352**. The inference engine **406** may process the received data to identify the sensitivity level of content from which the facets and/or metadata tags were derived. The inference engine **406** may be an artificial intelligence (AI) engine capable of processing natural language inputs such as, without limitation, the IBM Watson system. When the data is received at the inference engine **406**, the inference engine **406** automatically utilizes an artificial intelligence (AI) engine (e.g. IBM Watson, etc.) to perform advanced language processing on unstructured information using, for example but not limited to, the UIMA, using pre-existing knowledge base **400**, resulting in an identification of sensitivity level that corresponds to the content.

The inference engine **406** may automatically define the sensitivity levels and their correlation with different facets and/or metadata tags (e.g., UIMA) based on, for example, business needs, types of sensitive information encountered, potential uses and/or misuses of the sensitive information, available computing resources, or the like. Alternatively

and/or additionally, a user may define the sensitivity levels and their correlation with different facets and/or metadata tags. For example, a rule may associate a particular sensitivity level for content that includes PII. In the above example, the rules may specify different sensitivity levels for different types of PII included in the content and the corresponding security actions. For example, a lower sensitivity level may be associated with content if it includes names and date of births but not social security number(s). However, if a social security number(s) (SSN) is included in content, a higher sensitivity level may be assigned to the content.

The data handling agent **408** may then perform a security action on the content based on the identified sensitivity level and the rules associated with the identified sensitivity level (from the policies engine **404**). For example, considering the above PII examples, a rule may associate a sensitivity level for content that includes PII, and provide a security action such as a redaction security action to redact all PII before dissemination of the content. The rule may also specify one or more users or designations (e.g., HR) who may have access to content based on the sensitivity level (e.g., a first group of users may access the content after the redaction action security action is performed, a second group of users may access the content including PII without redaction action). As such, a security action may be performed to associate the appropriate access levels with the content. In the above example, the rules may specify different sensitivity levels for different types of PII included in the content and the corresponding security actions. For example, a lower sensitivity level may be associated with content if it includes names and date of births but not social security number(s), and the associated security action may specify that only redaction action should be performed. However, if a social security number(s) (SSN) is included in content, a higher sensitivity level may be assigned to the content and an associated security action may include encryption with distribution of keys only to certain users. Alternatively, security actions may specify different types of encryption based on the sensitivity level. For example, the level of encryption may be increased for more sensitive data compared to less sensitive data (e.g., a more secure algorithm/key-strength).

It will be understood to those skilled in the art that while this disclosure describes identification of a sensitivity level before associating a security action to be performed for the content, the principles disclosed herein may be used to directly associate various facets and/or metadata tags with security actions to be performed (without first identifying the sensitivity level). For example, if the facets include PII, then certain security actions may be performed relating to PII (and/or based on the type of PII) without first defining a sensitivity level. In another example, if facets and/or metadata tags associated with a content include the words “transaction”, “MasterCard”, and credit card number format, an appropriate security action may be performed relating to financial data.

The security module **360** may identify the sensitivity level before and/or after insertion of the metadata in the metadata store. For example, the security module **360** may analyze the metadata, the extracted facets and/or the tags to identify the sensitivity level upon receipt of metadata by the metadata collection system **351** before it is inserted into the metadata store **352**, and add the sensitivity level and/or security action information to the entry corresponding collected metadata in the metadata store **352**.

In certain embodiments, the security module **360** may perform the security action automatically after extraction of

the facets, upon tagging of the metadata and/or upon identification of a security level associated with the data (for example, when a user instruction is received to perform an intelligent protection of data in one or more data storage systems). The user instruction may include a query relating to tags and/or facets to identify the data to be protected by automatically identifying the appropriate security action. Alternatively, the user instruction may be a general instruction to intelligently protect all data residing in one or more data storage systems.

Additionally and/or alternatively, the security module **360** may perform the security action upon receipt of user instructions that may include, for example, instructions to perform a security action by defining a query relating to tags, facets and/or security level (e.g., encrypt data that includes personal information), or the like. If such a user instruction is received, the security module **360** does not need to scan all the data residing on the data storage systems to determine which data satisfies the user's query. Rather, the system may analyze the metadata stored in the metadata store, the facets, the tags, and/or the security levels to identify which data in the data storage systems satisfies the user's query. The user may also provide the security action to be performed once the data is identified. The security module **360** may then perform the requested security action on the identified data.

In one or more embodiments, the security module **360** may also include an alert miner **510** configured to extract indications of any new security breaches corresponding to one or more security actions performed by the security module **360** on the content residing on the data storage systems **320a**, **320b** . . . **320n** from one or more data sources (e.g., from text, images, video, audio, or the like). Examples of data sources may include new feeds, blogs, emails, social networking sites, or the like. For example, the alert miner **410** may subscribe to data mining systems such as the IBM Watson™ system (DISCOVERY NEWS) for receiving real-time alerts relating to any new software bugs, malware attacks, viruses, firewall breaches, or other types of data security breaches. In response to receiving an alert, the security module **360** may analyze if any content in the data storage systems **320a**, **320b** . . . **320n** may be affected (for example, by analyzing the security actions performed on data), and may then perform an appropriate remedial action. For example, if some content on the data storage systems **320a**, **320b** . . . **320n** is encrypted using an encryption technology that is subject to a malware attack, the security module **360** may update the encryption of the content to prevent any security breaches. Similarly, if the security module **360** receives an alert that content relating to an enterprise whose data is stored on the data storage systems **320a**, **320b** . . . **320n** has been compromised, the security module **360** may take appropriate action to install a firewall such that other content is not compromised, encrypt the enterprise's content, update its encryption (e.g., change encryption keys), and/or the like.

Referring now to FIG. **5**, an exemplary flowchart in accordance with various embodiments illustrating and describing a method of intelligent encryption of data is described. While the method **500** is described for the sake of convenience and not with an intent of limiting the disclosure as comprising a series and/or a number of steps, it is to be understood that the process does not need to be performed as a series of steps and/or the steps do not need to be performed in the order shown and described with respect to FIG. **5** but the process may be integrated and/or one or more

steps may be performed together, simultaneously, or the steps may be performed in the order disclosed or in an alternate order.

At step **502**, the metadata management system may collect metadata from one or more data storage systems. The metadata may include system metadata, scan metadata, and/or event metadata. As discussed above, the metadata management system may collect the event metadata (i.e., metadata corresponding a data operation executed on a data storage system) upon occurrence of every new event and/or periodically. The metadata management system may collect the event metadata by configuring the data storage systems to send event metadata to the metadata collection system of the metadata management system and/or by monitoring various data operations executed on the data storage systems or the network. The metadata management system may collect the scan metadata by periodically scanning the content of the data storage systems.

At step **504**, the metadata management system may analyze the metadata to identify various metadata tags to be applied to the metadata. In one or more embodiments, the metadata management system may identify the tags based on user defined policies that include a collection of classification rules and corresponding metadata tags for classification of metadata. As discussed above the user may provide the policies via a user interface such as a REST-API. The metadata management system may analyze the metadata (e.g., event metadata) using the user defined policies to identify various tags to be applied to the metadata. For example, a policy may include the following classification rule and the corresponding tag information:

```
owner:bob AND tag1:U* AND size:[5000 TO 5000]
"tag 1": "myFirstTag",
"tag2": 25
```

The metadata management system may then analyze the received metadata for "owner:bob AND tag1:U\* AND size:[5000 TO 5000]", and if a match is found, may apply a tag1 as "myFirstTag" and a tag 2 as "25" to the metadata. In an embodiment, the tag may provide information about the content corresponding to the metadata.

In one or more embodiments, the policies for applying the metadata tags may correspond to specific business needs, security levels of associated content (based on, for example, author, included information, file system, or the like), or the like. For example, if the owner is a physician, the tags may include information corresponding to access levels of the physician's patient files.

In certain embodiments, the metadata management system may first extract facets from the collected metadata and then use the extracted facets to classify the metadata and apply metadata tags. For example, if the facets include a medical practitioner's name and test reports extracted from the metadata of a documents (e.g., from owner information and the file name information), the metadata management system may identify the corresponding document to include patient information, and may apply tags such as confidential, privileged, patient personal information (PII), medical, or the like. Additionally and/or alternatively, the facets themselves may be used as tags—for example, the medical practitioner's name may be used as a tag to identify all documents that the medical practitioner has created/modified.

In certain embodiments, for extracting facets, the metadata management system may create annotators (dictionary based, entity based, custom, or the like) for classifying metadata stored in the metadata management system. As discussed above, the annotators may be created based on the

classification requirements. For example, if classification is to be performed to differentiate content associated with different entities or authors, an annotator may be created that identifies the occurrence of various entity/author names in the metadata. The system may then extract certain components of the collected metadata. Examples of the extracted components may include, without limitation, file name, file system path, owner information, object name, object bucket/container information, file size, modification history, or any other components that may provide information indicative of content corresponding to the metadata. The metadata management system may then extract facets from the extracted components of the metadata by passing the components through a classifier (e.g., UIMA), as discussed above.

In certain embodiments, the system may analyze the received event metadata and/or scan metadata in first-in, first-out manner. However, other orders for analyzing the data are within the scope of this disclosure.

The metadata management system may also directly extract facets from the data residing on the data storage systems (by, for example, using information included in file headers). For example, as discussed above, the system may perform supervised learning, unsupervised learning and/or deep inspection methods to extract facets from the candidate data. Optionally, facet extraction from the data residing on the data storage systems itself may be performed only for data on candidate data of step 506.

At step 506, the metadata management system may install the metadata in a database with the identified metadata tag and/or extracted facets (e.g., as indexes) in an appropriate format. For example, if the database is in the form of a key-value pair noSQL database, the system may sort and save the metadata in the form of various key-value pairs as described above. Furthermore, the tag(s) and/or facet(s) may also be applied as key-value pair(s).

At step 508, the metadata management system may identify a candidate data set on which an appropriate security action should be performed, where the candidate data set resides on a data storage system. In certain embodiments, the metadata management system may identify the candidate data set based on a query received from a user (e.g., based on queries defined by a user such as protect all data tagged as confidential, protect all data that include facets such as financial information, or personal information, etc.). In yet another embodiment, the metadata management system may automatically identify the candidate data based on the metadata tags and/or facets in 504. For example, if the extracted facets include PII, the metadata management system may automatically determine that a security action may need to be performed.

The metadata management system may automatically identify the candidate data in response to a user instruction, periodically, etc. Additionally and/or alternatively, the metadata management system may identify the candidate data set automatically in response to new metadata being received by the metadata management system (e.g., event metadata, when one or more events being executed; scan metadata if a scan is performed, etc.). For example, in response to a data operation event (e.g., copy, create, delete, open, close, read, write, backup, etc.) on a data set, the metadata management system may identify that data set as the candidate data set if it satisfies one or more rules for a security action. As such, the security action may be performed based on live events or scans and in a real-time manner. In certain embodiments, the metadata management system may analyze all events being executed on data residing on the data storage system to

determine if a security action should be performed. For example, if a copy operation is performed on a data set, the metadata management system may identify it as the candidate data set if the destination location corresponding to a copy operation does not have appropriate security protections. As such, a security action (e.g., encryption, backup, etc.) will need to be performed to prevent loss of data. The security action may be identified based on the candidate data set, destination location properties, etc., as discussed below. Similarly, if new data is written on a data storage system, the metadata management system may determine whether a security action should be performed on the data and the rules associated with it.

At 510, the metadata management system identifies an appropriate security action. The metadata management system identifies the appropriate security action based on, without limitation, the metadata tags, the facets extracted from metadata, and/or the facets extracted from the candidate data, and using one or more rules to identify a corresponding security action, as discussed above. In certain embodiments, the metadata management system may first identify a sensitivity level of the data based on, without limitation, the metadata tags, the facets extracted from metadata, and/or the facets extracted from the candidate data, and then identify the appropriate security action based on the sensitivity level.

In one or more embodiments, the security action may include an encryption action, and the sensitivity level (or the metadata tags, the facets extracted from metadata, and/or the facets) may be used to identify the appropriate encryption methods, levels, protocols, strength, access levels, or the like, for the candidate data. For example, if the extracted facets include financial information such as credit card numbers, the encryption required may be robust and may identify employees who may have access to the encrypted data (e.g., via sharing of encryption keys). On the other hand, if the extracted facets include personnel information such as full names and date of births, the encryption protocols may not be as robust, and a different set of people may be provided access to the data (e.g., human resources of the appropriate enterprise). Similarly, if the content includes, for example, on-demand videos, and the extracted facets provide the broadcasting channel name, the encryption methods may use encryption methods and access control lists corresponding to the broadcasting channel.

In certain embodiments, additional security level metadata tags may be added to the metadata stored in the metadata management system, and may provide a sensitivity level associated with content corresponding to the metadata. For example, if the extracted facets include PII, a corresponding sensitivity level may be identified using the security level metadata tag.

In one or more embodiment, the metadata management system may execute the identified security action on the identified data set (512).

At 514, the metadata management system may monitor real-time data to identify real-time threats to security protocols such as breaches in encryption protocols, breaches in firewalls, malware attacks, software bugs, brute force attacks, viruses, or the like. If the metadata management system identifies any threat to a security protocol that has been used to perform a security action on data residing on the data storage systems, the metadata management system may perform an appropriate remedial action (516).

While the illustrative embodiments described above are preferably implemented in hardware, such as in units and circuitry of a processor, various aspects of the illustrative



embodiments may be implemented in software as well. For example, it will be understood that each block of the flowchart illustrations in FIG. 6, and combinations of blocks in the flowchart illustration, can be implemented by computer program instructions. These computer program instructions may be provided to a processor or other programmable data processing apparatus to produce a machine, such that the instructions which execute on the processor or other programmable data processing apparatus create means for implementing the functions specified in the flowchart block or blocks. These computer program instructions may also be stored in a computer-readable memory or storage medium that can direct a processor or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory or storage medium produce an article of manufacture including instruction means which implement the functions specified in the flowchart block or blocks.

Accordingly, blocks of the flowchart illustration support combinations of means for performing the specified functions, combinations of steps for performing the specified functions, and program instruction means for performing the specified functions. It will also be understood that each block of the flowchart illustration, and combinations of blocks in the flowchart illustration, can be implemented by special purpose hardware-based computer systems which perform the specified functions or steps, or by combinations of special purpose hardware and computer instructions.

One or more embodiments of the present disclosure may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present disclosure.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway

computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present disclosure may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present disclosure.

Aspects of the present disclosure are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other

device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

Moreover, a system according to various embodiments may include a processor and logic integrated with and/or executable by the processor, the logic being configured to perform one or more of the process steps recited herein. By integrated with, what is meant is that the processor has logic embedded therewith as hardware logic, such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), etc. By executable by the processor, what is meant is that the logic is hardware logic; software logic such as firmware, part of an operating system, part of an application program; etc., or some combination of hardware and software logic that is accessible by the processor and configured to cause the processor to perform some functionality upon execution by the processor. Software logic may be stored on local and/or remote memory of any memory type, as known in the art. Any processor known in the art may be used, such as a software processor module and/or a hardware processor such as an ASIC, a FPGA, a central processing unit (CPU), an integrated circuit (IC), a graphics processing unit (GPU), etc.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the embodiments of the present disclosure has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the embodiments in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the disclosure. The embodiments and examples were chosen and described in order to best explain the principles of the disclosure and the practical application, and to enable others of ordinary skill in the art to understand the disclosure for various embodiments with various modifications as are suited to the particular use contemplated.

The programs described herein are identified based upon the application for which they are implemented in a specific embodiment of the disclosure. However, it should be appreciated that any particular program nomenclature herein is used merely for convenience, and thus the disclosure should not be limited to use solely in any specific application identified and/or implied by such nomenclature.

It will be clear that the various features of the foregoing systems and/or methodologies may be combined in any way, creating a plurality of combinations from the descriptions presented above.

It will be further appreciated that embodiments of the present disclosure may be provided in the form of a service deployed on behalf of a customer to offer service on demand.

The descriptions of the various embodiments of the present disclosure have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

What is claimed is:

1. A method of performing a security action on electronic data residing on an electronic data storage system to prevent subsequent unauthorized access comprising:

- maintaining an external metadata management system remote, separate, and independent from and in communication with a plurality of electronic data storage systems, wherein each electronic data storage system stores electronic data and comprises one or more processors having circuitry and logic for performing calculations and logic operations and wherein the external metadata management system comprises one or more metadata processors having circuitry and logic for performing calculations and logic operations;
- operating the external metadata management system by the one or more metadata processors, wherein operating the external metadata management system comprises storing metadata on the external metadata management system that corresponds to the electronic data residing on the plurality of data storage systems;
- identifying, by the external metadata management system using information included in the external metadata management system, a candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed, wherein the at least one of the plurality of data storage systems is remote, separate and independent from the external metadata management system and the at least one security action comprises at least one of a group consisting of: encryption; redaction; application of access controls; application of dissemination controls; inclusion and/or exclusion of partitions, folders, files, file extensions or other divisions of information; specifying encryption protocols and/or levels for encryption; applying and/or configuring operating system protection rules; applying and/or configuring user authentication mechanisms; applying and/or configuring device usage controls; auditing of system and/or user behavior; applying and/or configuring firewall rules; applying and/or configuring rules regarding anti-virus, anti-spam, and/or anti-spyware; applying and/or configuring rules regarding provisioning, management, and/or integration of third-party applications; and combinations thereof;
- in response to identifying the candidate data set residing on the at least one of the plurality of data storage systems on which at least one security action should be

performed, identifying by the external metadata management system the at least one security action that should be performed on the identified candidate data set; and

implementing, in response to identifying the at least one security action, the at least one identified security action on the identified candidate data set residing on the at least one of the plurality of data storage systems to prevent subsequent unauthorized access to the identified candidate data set residing on the at least one of the plurality of data storage systems.

2. The method of claim 1, wherein the at least one security action comprises encryption of the candidate data set.

3. The method of claim 1, wherein identifying the at least one security action comprises:

extracting one or more facets of the candidate data set stored with the metadata in the external metadata management system; and

using the one or more facets to identify the at least one security action.

4. The method of claim 3, wherein the extracted one or more facets are identified by performing data analytics on the candidate data set.

5. The method of claim 3, wherein the extracted one or more facets are identified by performing data analytics on at least one component of metadata corresponding to the candidate data set.

6. The method of claim 1, wherein identifying the at least one security action comprises:

identifying one or more custom tags for metadata corresponding to the candidate data set; and

using the one or more custom tags to identify the at least one security action.

7. The method of claim 1, wherein identifying the candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed comprises receiving a query from a user that includes one or more rules for selecting the candidate data set using metadata stored in the external metadata management system.

8. The method of claim 1, wherein identifying the candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed comprises identifying the candidate data set based on metadata received in response to a data operation event performed on the candidate data set.

9. The method of claim 1, further comprising:

identifying a security level associated with the candidate data set based on at least one of a group consisting of: one or more facets extracted from the candidate data set and stored with the metadata in the external metadata management system, one or more facets extracted from metadata associated with the candidate data set and stored with the metadata in the external metadata management system, one or more custom tags corresponding to metadata associated with the candidate data set, and combinations thereof; and

using the security level to identify the at least one security action.

10. The method of claim 1, further comprising:

receiving a real-time alert comprising a threat to the at least one security action; and

identifying at least one remedial action for countering the threat.

11. A non-transitory computer readable medium comprising programming instructions that when executed cause one or more metadata processors to:

maintain an external metadata management system remote, separate, and independent from and in communication with a plurality of data storage systems, wherein each electronic data storage system stores electronic data and comprises one or more processors having circuitry and logic for performing calculations and logic operations and wherein the external metadata management system comprises the one or more metadata processors where the one or more metadata processors have circuitry and logic for performing calculations and logic operations;

store metadata on the external metadata management system that corresponds to the electronic data residing on the plurality of data storage systems;

identify, by the external metadata management system using information included in the external metadata management system, a candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed, wherein the at least one of the plurality of data storage systems is remote, separate, and independent of the external metadata management system and the at least one security action comprises at least one of a group consisting of: encryption; redaction; application of access controls; application of dissemination controls; inclusion and/or exclusion of partitions, folders, files, file extensions or other divisions of information; specifying encryption protocols and/or levels for encryption; applying and/or configuring operating system protection rules; applying and/or configuring user authentication mechanisms; applying and/or configuring device usage controls; auditing of system and/or user behavior; applying and/or configuring firewall rules; applying and/or configuring rules regarding anti-virus, anti-spam, and/or anti-spyware; applying and/or configuring rules regarding provisioning, management, and/or integration of third-party applications; and combinations thereof;

in response to identifying the candidate data set residing on at least one of the plurality of data storage systems on which at least one security action should be performed, identify by the external metadata management system the at least one security action; and

implement, in response to identifying the at least one security action, the at least one identified security action on the identified candidate data set residing on the at least one of the plurality of data storage systems to prevent subsequent unauthorized access to the identified candidate data set residing on the at least one of the plurality of data storage systems.

12. The non-transitory computer readable medium of claim 11, wherein the at least one security action comprises encryption of the candidate data set.

13. The non-transitory computer readable medium of claim 11, wherein causing the one or more metadata processors to identify the at least one security action comprises causing the processor to:

extract one or more facets of the candidate data set stored with the metadata in the external metadata management system; and

use the one or more facets to identify the at least one security action.

14. The non-transitory computer readable medium of claim 13, wherein the extracted one or more facets are identified by performing data analytics on the candidate data set.

39

15. The non-transitory computer readable medium of claim 13, wherein the extracted one or more facets are identified by performing data analytics on at least one component of metadata corresponding to the candidate data set.

16. The non-transitory computer readable medium of claim 11, wherein causing the one or more metadata processors to identify the at least one security action comprises causing the processor to:

identify one or more custom tags for metadata corresponding to the candidate data set; and

use the one or more custom tags to identify the at least one security action.

17. The non-transitory computer readable medium of claim 11, further comprising programming instruction that when executed cause the one or more metadata processors to:

identify a security level associated with the candidate data set based on at least one of a group consisting of: one

40

or more facets extracted from the candidate data set and stored with the metadata in the external metadata management system, one or more facets extracted from metadata associated with the candidate data set and stored with the metadata in the external metadata management system, one or more custom tags corresponding to metadata associated with the candidate data set, and combinations thereof; and use the security level to identify the at least one security action.

18. The non-transitory computer readable medium of claim 11, further comprising programming instruction that when executed cause the one or more metadata processors to:

receive a real-time alert comprising a threat to the at least one security action; and

identify at least one remedial action for countering the threat.

\* \* \* \* \*